



(12) **United States Patent**
Kottomtharayil et al.

(10) **Patent No.:** **US 11,734,035 B2**
(45) **Date of Patent:** ***Aug. 22, 2023**

(54) **VIRTUAL MACHINE LOAD BALANCING**

(71) Applicant: **Commvault Systems, Inc.**, Tinton Falls, NJ (US)

(72) Inventors: **Rajiv Kottomtharayil**, Marlboro, NJ (US); **Rahul S. Pawar**, Marlboro, NJ (US); **Ashwin Gautamchand Sancheti**, Marlboro, NJ (US); **Sumer Dilip Deshpande**, Ocean, NJ (US); **Sri Karthik Bhagi**, Morganville, NJ (US); **Henry Wallace Dornemann**, Eatontown, NJ (US); **Ananda Venkatesha**, Cupertino, CA (US)

(73) Assignee: **Commvault Systems, Inc.**, Tinton Falls, NJ (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 138 days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **17/118,329**

(22) Filed: **Dec. 10, 2020**

(65) **Prior Publication Data**

US 2021/0096893 A1 Apr. 1, 2021

Related U.S. Application Data

(63) Continuation of application No. 16/518,107, filed on Jul. 22, 2019, now Pat. No. 10,896,053, which is a (Continued)

(51) **Int. Cl.**
G06F 9/455 (2018.01)
G06F 9/50 (2006.01)

(52) **U.S. Cl.**
CPC **G06F 9/45533** (2013.01); **G06F 9/505** (2013.01); **G06F 9/5083** (2013.01); **G06F 2209/5022** (2013.01)

(58) **Field of Classification Search**

CPC G06F 9/45533
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,084,231 A 4/1978 Capozzi et al.
4,267,568 A 5/1981 Dechant et al.
(Continued)

FOREIGN PATENT DOCUMENTS

EP 0259912 A1 3/1988
EP 0405926 A2 1/1991
(Continued)

OTHER PUBLICATIONS

U.S. Appl. No. 16/262,753, filed Jan. 30, 2019, Dornemann et al.
(Continued)


Primary Examiner — Charlie Sun

(74) *Attorney, Agent, or Firm* — Commvault Systems, Inc.

(57) **ABSTRACT**

Virtual machine (VM) proliferation may be reduced by determining the availability of existing VMs to perform a task. Tasks may be assigned to existing VMs instead of creating a new VM to perform the task. Furthermore, a coordinator may determine a grouping of VMs or VM hosts based on one or more factors associated with the VMs or the VM hosts, such as VM type or geographical location of the VM hosts. The coordinator may also assign one or more Virtual Server Agents (VSAs) to facilitate managing the group of VM hosts. In some embodiments, the coordinators may facilitate load balancing of VSAs during operation, such as during a backup operation, a restore operation, or any other operation between a primary storage system and a secondary storage system.

22 Claims, 11 Drawing Sheets



Related U.S. Application Data

continuation of application No. 15/960,401, filed on Apr. 23, 2018, now Pat. No. 10,474,483, which is a continuation of application No. 15/196,739, filed on Jun. 29, 2016, now Pat. No. 9,977,687, which is a continuation of application No. 14/148,507, filed on Jan. 6, 2014, now Pat. No. 9,703,584.

(60) Provisional application No. 61/750,255, filed on Jan. 8, 2013.

(56) References Cited

U.S. PATENT DOCUMENTS

4,283,787	A	8/1981	Chambers	6,131,095	A	10/2000	Low et al.
4,417,321	A	11/1983	Chang et al.	6,131,190	A	10/2000	Sidwell
4,641,274	A	2/1987	Swank	6,148,412	A	11/2000	Cannon et al.
4,654,819	A	3/1987	Stiffler et al.	6,154,787	A	11/2000	Urevig et al.
4,686,620	A	8/1987	Ng	6,161,111	A	12/2000	Mutalik et al.
4,912,637	A	3/1990	Sheedy et al.	6,167,402	A	12/2000	Yeager
4,995,035	A	2/1991	Cole et al.	6,212,512	B1	4/2001	Barney et al.
5,005,122	A	4/1991	Griffin et al.	6,260,069	B1	7/2001	Anglin
5,093,912	A	3/1992	Dong et al.	6,269,431	B1	7/2001	Dunham
5,133,065	A	7/1992	Cheffetz et al.	6,275,953	B1	8/2001	Vahalia et al.
5,193,154	A	3/1993	Kitajima et al.	6,301,592	B1	10/2001	Aoyama et al.
5,212,772	A	5/1993	Masters	6,324,581	B1	11/2001	Xu et al.
5,226,157	A	7/1993	Nakano et al.	6,328,766	B1	12/2001	Long
5,239,647	A	8/1993	Anglin et al.	6,330,570	B1	12/2001	Crighton
5,241,668	A	8/1993	Eastridge et al.	6,330,642	B1	12/2001	Carteau
5,241,670	A	8/1993	Eastridge et al.	6,343,324	B1	1/2002	Hubis et al.
5,276,860	A	1/1994	Fortier et al.	RE37,601	E	3/2002	Eastridge et al.
5,276,867	A	1/1994	Kenley et al.	6,356,801	B1	3/2002	Goodman et al.
5,287,500	A	2/1994	Stoppioni, Jr.	6,389,432	B1	5/2002	Pothapragada et al.
5,301,286	A	4/1994	Rajani	6,397,242	B1	5/2002	Devine et al.
5,321,816	A	6/1994	Rogan et al.	6,418,478	B1	7/2002	Ignatius et al.
5,333,315	A	7/1994	Saether et al.	6,421,711	B1	7/2002	Blumenau et al.
5,347,653	A	9/1994	Flynn et al.	6,487,561	B1	11/2002	Ofek et al.
5,410,700	A	4/1995	Fecteau et al.	6,519,679	B2	2/2003	Devireddy et al.
5,420,996	A	5/1995	Aoyagi	6,538,669	B1	3/2003	Lagueux, Jr.
5,448,724	A	9/1995	Hayashi	6,542,972	B2	4/2003	Ignatius et al.
5,454,099	A	9/1995	Myers et al.	6,564,228	B1	5/2003	O'Connor
5,491,810	A	2/1996	Allen	6,581,076	B1	6/2003	Ching et al.
5,495,607	A	2/1996	Pisello et al.	6,658,436	B2	12/2003	Oshinsky et al.
5,504,873	A	4/1996	Martin et al.	6,658,526	B2	12/2003	Nguyen et al.
5,544,345	A	8/1996	Carpenter et al.	6,721,767	B2	4/2004	De Meno et al.
5,544,347	A	8/1996	Yanai et al.	6,760,723	B2	7/2004	Oshinsky et al.
5,559,957	A	9/1996	Balk	6,772,290	B1	8/2004	Bromley et al.
5,559,991	A	9/1996	Kanfi	6,820,214	B1	11/2004	Cabrera et al.
5,594,901	A	1/1997	Andoh	7,003,641	B2	2/2006	Prahlad et al.
5,619,644	A	4/1997	Crockett et al.	7,035,880	B1	4/2006	Crescenti et al.
5,638,509	A	6/1997	Dunphy et al.	7,076,270	B2	7/2006	Jaggers et al.
5,642,496	A	6/1997	Kanfi	7,107,298	B2	9/2006	Prahlad et al.
5,664,204	A	9/1997	Wang	7,130,970	B2	10/2006	Devassy et al.
5,673,381	A	9/1997	Huai et al.	7,162,496	B2	1/2007	Amarendran et al.
5,699,361	A	12/1997	Ding et al.	7,174,433	B2	2/2007	Kottomtharayil et al.
5,729,743	A	3/1998	Squibb	7,219,162	B2	5/2007	Donker et al.
5,751,997	A	5/1998	Kullick et al.	7,246,207	B2	7/2007	Kottomtharayil
5,758,359	A	5/1998	Saxon	7,315,923	B2	1/2008	Retnamma et al.
5,761,677	A	6/1998	Senator et al.	7,324,543	B2	1/2008	Wassew et al.
5,764,972	A	6/1998	Crouse et al.	7,343,356	B2	3/2008	Prahlad et al.
5,778,395	A	7/1998	Whiting et al.	7,343,453	B2	3/2008	Prahlad et al.
5,812,398	A	9/1998	Nielsen	7,346,751	B2	3/2008	Prahlad et al.
5,813,009	A	9/1998	Johnson et al.	7,356,679	B1	4/2008	Le
5,813,017	A	9/1998	Morris	7,376,895	B2	5/2008	Tsao
5,875,478	A	2/1999	Blumenau	7,386,744	B2	6/2008	Barr et al.
5,887,134	A	3/1999	Ebrahim	7,389,311	B1	6/2008	Crescenti et al.
5,901,327	A	5/1999	Ofek	7,395,282	B1	7/2008	Crescenti
5,924,102	A	7/1999	Perks	7,440,982	B2	10/2008	Lu et al.
5,950,205	A	9/1999	Aviani, Jr.	7,448,079	B2	11/2008	Tremain
5,974,563	A	10/1999	Beeler, Jr.	7,454,569	B2	11/2008	Kavuri et al.
6,021,415	A	2/2000	Cannon et al.	7,475,282	B2	1/2009	Tormasov et al.
6,026,414	A	2/2000	Anglin	7,484,208	B1	1/2009	Nelson
6,052,735	A	4/2000	Ulrich et al.	7,490,207	B2	2/2009	Amarendran et al.
6,076,148	A	6/2000	Kedem	7,500,053	B1	3/2009	Kavuri et al.
6,094,416	A	7/2000	Ying	7,502,820	B2	3/2009	Manders et al.
6,101,585	A	8/2000	Brown et al.	7,529,782	B2	5/2009	Prahlad et al.
				7,536,291	B1	5/2009	Vijayan Retnamma et al.
				7,543,125	B2	6/2009	Gokhale
				7,546,324	B2	6/2009	Prahlad et al.
				7,552,279	B1	6/2009	Gandler
				7,599,971	B1	10/2009	Hagerstrom
				7,603,386	B2	10/2009	Amarendran et al.
				7,606,844	B2	10/2009	Kottomtharayil
				7,613,752	B2	11/2009	Prahlad et al.
				7,617,253	B2	11/2009	Prahlad et al.
				7,617,262	B2	11/2009	Prahlad et al.
				7,620,710	B2	11/2009	Kottomtharayil et al.
				7,631,351	B2	12/2009	Erofeev
				7,636,743	B2	12/2009	Erofeev
				7,640,406	B1	12/2009	Hagerstrom et al.
				7,651,593	B2	1/2010	Prahlad et al.
				7,657,550	B2	2/2010	Prahlad et al.
				7,660,807	B2	2/2010	Prahlad et al.

(56)

References Cited

U.S. PATENT DOCUMENTS

7,661,028 B2	2/2010	Erofeev	8,577,845 B2	11/2013	Nguyen et al.
7,668,884 B2	2/2010	Prahlad et al.	8,578,120 B2	11/2013	Attarde et al.
7,685,177 B1	3/2010	Hagerstrom et al.	8,578,126 B1	11/2013	Gaonkar et al.
7,694,070 B2	4/2010	Mogi et al.	8,606,911 B2	12/2013	Raleigh et al.
7,716,171 B2	5/2010	Kryger	8,612,439 B2	12/2013	Prahlad et al.
7,721,138 B1	5/2010	Lyadvinsky et al.	8,615,489 B2	12/2013	Pershin
7,725,671 B2	5/2010	Prahlad et al.	8,620,870 B2	12/2013	Dwarampudi et al.
7,725,893 B2	5/2010	Jaeckel et al.	8,621,460 B2	12/2013	Evans et al.
7,730,035 B2	6/2010	Berger et al.	8,635,429 B1	1/2014	Naftel et al.
7,734,669 B2	6/2010	Kottomtharayil et al.	8,667,171 B2	3/2014	Guo et al.
7,747,579 B2	6/2010	Prahlad et al.	8,677,085 B2	3/2014	Vaghani et al.
7,756,835 B2	7/2010	Pugh	8,719,286 B1	5/2014	Xing
7,756,964 B2	7/2010	Madison, Jr. et al.	8,751,857 B2	6/2014	Frenkel et al.
7,765,167 B2	7/2010	Prahlad et al.	8,776,043 B1	7/2014	Thimmes et al.
7,778,984 B2	8/2010	Zhang et al.	8,799,431 B2	8/2014	Pabari
7,792,789 B2	9/2010	Prahlad et al.	8,831,202 B1	9/2014	Abidogun et al.
7,793,307 B2	9/2010	Gokhale et al.	8,844,015 B2	9/2014	Pendergrass et al.
7,801,864 B2	9/2010	Prahlad et al.	8,850,146 B1	9/2014	Majumdar
7,809,914 B2	10/2010	Kottomtharayil et al.	8,904,081 B1	12/2014	Kulkarni
7,822,967 B2	10/2010	Fung	8,924,967 B2	12/2014	Nelson
7,823,145 B1	10/2010	Le et al.	8,930,543 B2	1/2015	Ashok et al.
7,831,789 B1	11/2010	Per et al.	8,938,481 B2	1/2015	Kumarasamy et al.
7,840,537 B2	11/2010	Gokhale	8,938,643 B1	1/2015	Karmarkar et al.
7,861,234 B1	12/2010	Lolo	8,954,446 B2	2/2015	Vijayan Retnamma et al.
7,882,077 B2	2/2011	Gokhale et al.	8,954,796 B1	2/2015	Cohen et al.
7,890,467 B2	2/2011	Watanable et al.	8,959,509 B1	2/2015	Sobel et al.
7,899,788 B2	3/2011	Chandhok et al.	8,966,318 B1	2/2015	Shah
7,917,617 B1	3/2011	Ponnapur	9,020,895 B1	4/2015	Rajashekar
7,925,850 B1	4/2011	Waldspurger et al.	9,020,900 B2	4/2015	Vijayan Retnamma et al.
7,937,421 B2	5/2011	Mikesell et al.	9,021,459 B1	4/2015	Qu
7,937,612 B1	5/2011	Lyadvinsky et al.	9,026,498 B2	5/2015	Kumarasamy
8,001,277 B2	8/2011	Mega et al.	9,069,587 B2	6/2015	Agarwal et al.
8,037,016 B2	10/2011	Odulinski et al.	9,098,457 B2	8/2015	Towstopiat et al.
8,037,028 B2	10/2011	Prahlad et al.	9,098,495 B2	8/2015	Gokhale
8,037,032 B2	10/2011	Pershin et al.	9,116,633 B2	8/2015	Sancheti et al.
8,046,550 B2	10/2011	Feathergill	9,124,633 B1	9/2015	Eizadi et al.
8,055,745 B2	11/2011	Atluri	9,141,529 B2	9/2015	Klein et al.
8,060,476 B1	11/2011	Afonso et al.	9,146,755 B2	9/2015	Lassonde et al.
8,069,271 B2	11/2011	Brunet et al.	9,213,706 B2	12/2015	Long et al.
8,099,391 B1	1/2012	Monckton	9,223,596 B1	12/2015	Araujo
8,117,492 B1	2/2012	Searls et al.	9,223,597 B2	12/2015	Deshpande et al.
8,135,930 B1	3/2012	Mattox et al.	9,235,474 B1	1/2016	Petri et al.
8,140,786 B2	3/2012	Bunte	9,235,582 B1	1/2016	Madiraju Varadaraju et al.
8,156,301 B1	4/2012	Khandelwal et al.	9,239,687 B2	1/2016	Vijayan et al.
8,170,995 B2	5/2012	Prahlad et al.	9,239,762 B1	1/2016	Gunda et al.
8,180,735 B2	5/2012	Ansari	9,246,996 B1	1/2016	Brooker
8,185,893 B2	5/2012	Hyser et al.	9,268,602 B2	2/2016	Prahlad et al.
8,191,063 B2	5/2012	Shingai et al.	9,280,378 B2	3/2016	Shah
8,200,637 B1	6/2012	Stringham	9,286,086 B2	3/2016	Deshpande et al.
8,200,638 B1	6/2012	Zheng	9,286,110 B2	3/2016	Mitkar et al.
8,209,680 B1	6/2012	Le et al.	9,292,350 B1	3/2016	Pendharkar et al.
8,219,524 B2	7/2012	Gokhale	9,298,715 B2	3/2016	Kumarasamy et al.
8,219,653 B1	7/2012	Keagy et al.	9,311,121 B2	4/2016	Deshpande et al.
8,219,769 B1	7/2012	Wilk	9,311,248 B2	4/2016	Wagner
8,225,133 B1	7/2012	Lyadvinsky et al.	9,311,313 B2	4/2016	Le et al.
8,229,896 B1	7/2012	Narayanan	9,354,927 B2	5/2016	Hiltgen et al.
8,229,954 B2	7/2012	Kottomtharayil et al.	9,397,944 B1	7/2016	Hobbs et al.
8,230,256 B1	7/2012	Raut	9,405,763 B2	8/2016	Prahlad et al.
8,285,681 B2	10/2012	Prahlad et al.	9,417,968 B2	8/2016	Dornemann et al.
8,307,177 B2	11/2012	Prahlad et al.	9,424,136 B1	8/2016	Teater et al.
8,307,187 B2	11/2012	Chawla et al.	9,436,555 B2	9/2016	Dornemann et al.
8,315,992 B1	11/2012	Gipp et al.	9,451,023 B2	9/2016	Sancheti
8,364,652 B2	1/2013	Vijayan et al.	9,461,881 B2	10/2016	Kumarasamy
8,370,542 B2	2/2013	Lu et al.	9,471,441 B1	10/2016	Lyadvinsky et al.
8,387,045 B2	2/2013	Yasutaka et al.	9,477,683 B2	10/2016	Ghosh
8,396,838 B2	3/2013	Brockway	9,489,244 B2	11/2016	Mitkar et al.
8,407,190 B2	3/2013	Prahlad	9,495,370 B1	11/2016	Chatterjee et al.
8,433,679 B2	4/2013	Crescenti	9,495,404 B2	11/2016	Kumarasamy et al.
8,434,131 B2	4/2013	Varadharajan	9,535,798 B1	1/2017	Liguori
8,438,347 B1	5/2013	Tawri et al.	9,575,789 B1	2/2017	Rangari et al.
8,453,145 B1	5/2013	Naik	9,575,991 B2	2/2017	Ghosh
8,458,419 B2	6/2013	Basler et al.	9,588,847 B1	3/2017	Natanzon et al.
8,473,594 B2	6/2013	Astete et al.	9,594,636 B2	3/2017	Mortensen et al.
8,489,676 B1	7/2013	Chaplin et al.	9,606,745 B2	3/2017	Satoyama et al.
8,560,788 B1	10/2013	Sreedharan et al.	9,612,966 B2	4/2017	Joshi et al.
			9,613,064 B1	4/2017	Chou et al.
			9,632,882 B2	4/2017	Kumarasamy et al.
			9,633,033 B2	4/2017	Vijayan et al.
			9,639,274 B2	5/2017	Maranna et al.

(56)

References Cited

U.S. PATENT DOCUMENTS

9,639,426	B2	5/2017	Pawar et al.	2005/0080970	A1	4/2005	Jeyasingh et al.
9,639,428	B1	5/2017	Boda et al.	2005/0198303	A1	9/2005	Knauerhase et al.
9,652,283	B2	5/2017	Mitkar et al.	2005/0216788	A1	9/2005	Mani-Meitav et al.
9,665,386	B2	5/2017	Bayapuneni et al.	2005/0262097	A1	11/2005	Sim-Tang
9,684,535	B2	6/2017	Deshpande et al.	2006/0058994	A1	3/2006	Ravi et al.
9,684,567	B2	6/2017	Derk et al.	2006/0064555	A1	3/2006	Prahlad et al.
9,703,584	B2	7/2017	Kottomtharayil et al.	2006/0101189	A1	5/2006	Chandrasekaran et al.
9,710,465	B2	7/2017	Dornemann et al.	2006/0155712	A1	7/2006	Prahlad et al.
9,766,989	B2	7/2017	Mitkar et al.	2006/0184935	A1	8/2006	Abels et al.
9,740,702	B2	8/2017	Pawar et al.	2006/0195715	A1	8/2006	Herington
9,740,723	B2	8/2017	Prahlad et al.	2006/0224846	A1	10/2006	Amarendran
9,760,398	B1	9/2017	Pai	2006/0225065	A1	10/2006	Chandhok et al.
9,760,448	B1	9/2017	Per et al.	2006/0230136	A1	10/2006	Ma
9,766,825	B2	9/2017	Bhagi et al.	2006/0259908	A1	11/2006	Bayer
9,792,075	B1	10/2017	Banerjee et al.	2007/0027999	A1	2/2007	Allen et al.
9,823,977	B2	11/2017	Dornemann et al.	2007/0043870	A1	2/2007	Ninose
9,852,026	B2	12/2017	Mitkar et al.	2007/0100792	A1	5/2007	Lent et al.
9,928,001	B2	3/2018	Dornemann et al.	2007/0198802	A1	8/2007	Kavuri
9,939,981	B2	4/2018	White et al.	2007/0203938	A1	8/2007	Prahlad et al.
9,965,316	B2	5/2018	Deshpande et al.	2007/0208918	A1	9/2007	Harbin et al.
9,977,687	B2	5/2018	Kottomtharayil et al.	2007/0220319	A1	9/2007	Desai et al.
9,983,936	B2	5/2018	Dornemann et al.	2007/0234302	A1	10/2007	Suzuki et al.
9,996,287	B2	6/2018	Dornemann et al.	2007/0239804	A1	10/2007	Armstrong et al.
9,996,429	B1	6/2018	Kumar	2007/0260831	A1	11/2007	Michael et al.
9,996,534	B2	6/2018	Dornemann et al.	2007/0266056	A1	11/2007	Stacey et al.
10,048,889	B2	8/2018	Dornemann et al.	2007/0288536	A1	12/2007	Sen et al.
10,061,658	B2	8/2018	Long et al.	2007/0300220	A1	12/2007	Seliger
10,075,459	B1	9/2018	Suryanarayanan	2008/0007765	A1	1/2008	Ogata et al.
10,108,652	B2	10/2018	Kumarasamy et al.	2008/0059704	A1	3/2008	Kavuri
10,114,705	B1	10/2018	Kumar et al.	2008/0071841	A1	3/2008	Okada et al.
10,152,251	B2	12/2018	Sancheti et al.	2008/0091655	A1	4/2008	Gokhale
10,162,528	B2	12/2018	Sancheti et al.	2008/0126833	A1	5/2008	Callaway et al.
10,162,873	B2	12/2018	Desphande et al.	2008/0134177	A1	6/2008	Fitzgerald et al.
10,210,048	B2	2/2019	Sancheti	2008/0189468	A1	8/2008	Schmidt et al.
10,225,964	B2	3/2019	Smith	2008/0195639	A1	8/2008	Freeman et al.
10,228,962	B2	3/2019	Dornemann et al.	2008/0228771	A1	9/2008	Prahlad et al.
10,241,871	B1	3/2019	Cheng	2008/0228833	A1	9/2008	Kano
10,387,073	B2	8/2019	Bhagi et al.	2008/0229037	A1	9/2008	Bunte
10,417,102	B2	9/2019	Sanakkayala et al.	2008/0235479	A1	9/2008	Scales et al.
10,437,505	B2	10/2019	Dornemann et al.	2008/0243855	A1	10/2008	Prahlad
10,452,303	B2	10/2019	Dornemann et al.	2008/0243947	A1	10/2008	Kaneda
10,474,483	B2	11/2019	Kottomtharayil et al.	2008/0244028	A1	10/2008	Le et al.
10,474,542	B2	11/2019	Mitkar et al.	2008/0244068	A1	10/2008	Iyoda et al.
10,474,548	B2	11/2019	Sanakkayala et al.	2008/0244177	A1	10/2008	Crescenti et al.
10,481,984	B1	11/2019	Semyonov et al.	2008/0250407	A1	10/2008	Dadhia et al.
10,496,547	B1	12/2019	Naenko	2008/0270564	A1	10/2008	Rangegowda et al.
10,565,067	B2	2/2020	Dornemann	2008/0275924	A1	11/2008	Fries
10,567,067	B2	2/2020	Kumagai	2008/0282253	A1	11/2008	Huizenga
10,572,468	B2	2/2020	Dornemann et al.	2008/0313371	A1	12/2008	Kedem et al.
10,592,350	B2	3/2020	Dornemann	2008/0320319	A1	12/2008	Muller
10,650,057	B2	5/2020	Pawar et al.	2009/0006733	A1	1/2009	Gold et al.
10,678,758	B2	6/2020	Dornemann	2009/0037680	A1	2/2009	Colbert et al.
10,684,883	B2	6/2020	Deshpande et al.	2009/0113109	A1	4/2009	Nelson et al.
10,747,630	B2	8/2020	Sanakkayala et al.	2009/0144416	A1	6/2009	Chatley et al.
10,768,971	B2	9/2020	Dornemann et al.	2009/0157882	A1	6/2009	Kashyap
10,776,209	B2	9/2020	Pawar et al.	2009/0210427	A1	8/2009	Eidler et al.
10,877,851	B2	12/2020	Mitkar et al.	2009/0210458	A1	8/2009	Glover et al.
10,877,928	B2	12/2020	Nagrle et al.	2009/0210527	A1	8/2009	Kawato
10,896,053	B2	1/2021	Kottomtharayil et al.	2009/0216816	A1	8/2009	Basler et al.
10,896,100	B2	1/2021	Mitkar et al.	2009/0222496	A1	9/2009	Liu et al.
10,983,875	B2	4/2021	Mitkar et al.	2009/0228669	A1	9/2009	Siesarev et al.
11,010,011	B2	5/2021	Varadharajan et al.	2009/0240904	A1	9/2009	Austruy et al.
11,249,864	B2	2/2022	Bhagi et al.	2009/0248762	A1	10/2009	Prahlad et al.
11,422,709	B2	8/2022	Dornemann et al.	2009/0249005	A1	10/2009	Bender et al.
11,436,202	B2	9/2022	Dornemann	2009/0282404	A1	11/2009	Khandekar et al.
2002/0069369	A1	6/2002	Tremain	2009/0287665	A1	11/2009	Prahlad
2002/0095609	A1	7/2002	Tokunaga	2009/0300023	A1	12/2009	Vaghani
2002/0194511	A1	12/2002	Swoboda	2009/0300057	A1	12/2009	Friedman
2003/0031127	A1	2/2003	Saleh et al.	2009/0307166	A1	12/2009	Routray et al.
2003/0126494	A1	7/2003	Strasser	2009/0313260	A1	12/2009	Mimatsu
2003/0204597	A1	10/2003	Arakawa et al.	2009/0313447	A1	12/2009	Nguyen et al.
2004/0030668	A1	2/2004	Pawlowski et al.	2009/0313503	A1	12/2009	Atluri et al.
2004/0030822	A1	2/2004	Rajan et al.	2009/0319534	A1	12/2009	Gokhale
2004/0230899	A1	11/2004	Pagnano et al.	2009/0319585	A1	12/2009	Gokhale
2005/0060356	A1	3/2005	Saika	2009/0320029	A1	12/2009	Kottomtharayil
				2009/0320137	A1	12/2009	White et al.
				2009/0327477	A1	12/2009	Madison, Jr. et al.
				2010/0011178	A1	1/2010	Feathergill
				2010/0017647	A1	1/2010	Callaway et al.

(56)		References Cited	
U.S. PATENT DOCUMENTS			
2010/0030984	A1	2/2010	Erickson
2010/0049929	A1	2/2010	Nagarkar et al.
2010/0049930	A1	2/2010	Pershin
2010/0070466	A1	3/2010	Prahlad et al.
2010/0070474	A1	3/2010	Lad
2010/0070725	A1	3/2010	Prahlad et al.
2010/0070726	A1	3/2010	Ngo et al.
2010/0077165	A1	3/2010	Lu
2010/0082672	A1	4/2010	Kottomtharayil
2010/0094948	A1	4/2010	Ganesh et al.
2010/0106691	A1	4/2010	Preslan et al.
2010/0107158	A1	4/2010	Chen et al.
2010/0107172	A1	4/2010	Calinescu et al.
2010/0115332	A1	5/2010	Zheng et al.
2010/0161919	A1	6/2010	Dodgson et al.
2010/0186014	A1	7/2010	Vaghani et al.
2010/0211829	A1	8/2010	Ziskind et al.
2010/0228913	A1	9/2010	Czezatke et al.
2010/0242096	A1	9/2010	Varadharajan et al.
2010/0257523	A1	10/2010	Frank
2010/0262585	A1	10/2010	Rosikiewicz et al.
2010/0262586	A1	10/2010	Rosikiewicz et al.
2010/0262794	A1	10/2010	De Beer et al.
2010/0280999	A1	11/2010	Atluri et al.
2010/0299309	A1	11/2010	Maki et al.
2010/0299666	A1	11/2010	Agbaria et al.
2010/0306173	A1	12/2010	Frank
2010/0306486	A1	12/2010	Balasubramanian et al.
2010/0325471	A1	12/2010	Mishra et al.
2010/0325727	A1	12/2010	Neystad et al.
2010/0332401	A1	12/2010	Prahlad
2010/0332454	A1	12/2010	Prahlad et al.
2010/0332456	A1	12/2010	Prahlad et al.
2010/0332479	A1	12/2010	Prahlad
2010/0332629	A1	12/2010	Cotugno et al.
2010/0332818	A1	12/2010	Prahlad
2010/0333100	A1	12/2010	Miyazaki et al.
2010/0333116	A1	12/2010	Prahlad
2011/0004586	A1	1/2011	Cherryholmes et al.
2011/0010515	A1	1/2011	Ranade
2011/0016467	A1	1/2011	Kane
2011/0022811	A1	1/2011	Kirihata et al.
2011/0022812	A1	1/2011	Van Der Linden et al.
2011/0023114	A1	1/2011	Diab et al.
2011/0035620	A1	2/2011	Elyashev et al.
2011/0047541	A1	2/2011	Yamaguchi et al.
2011/0061045	A1	3/2011	Phillips
2011/0072430	A1	3/2011	Mani
2011/0087632	A1	4/2011	Subramanian et al.
2011/0107025	A1	5/2011	Urkude et al.
2011/0107331	A1	5/2011	Evans et al.
2011/0138069	A1	6/2011	Momchilov et al.
2011/0161299	A1	6/2011	Prahlad
2011/0179414	A1	7/2011	Goggin et al.
2011/0185355	A1	7/2011	Chawla et al.
2011/0191559	A1	8/2011	Li et al.
2011/0196840	A1	8/2011	Barzilai
2011/0202728	A1	8/2011	Nichols et al.
2011/0202734	A1	8/2011	Dhakras et al.
2011/0208928	A1	8/2011	Chandra et al.
2011/0213754	A1	9/2011	Bindal
2011/0219144	A1	9/2011	Amit et al.
2011/0225277	A1	9/2011	Freimuth et al.
2011/0239013	A1	9/2011	Muller
2011/0246430	A1	10/2011	Prahlad et al.
2011/0252208	A1	10/2011	Ali et al.
2011/0264786	A1	10/2011	Kedem et al.
2012/0002951	A1	1/2012	Reisman
2012/0016840	A1	1/2012	Lin et al.
2012/0017027	A1	1/2012	Baskakov et al.
2012/0017043	A1	1/2012	Aizman et al.
2012/0017114	A1	1/2012	Timashev et al.
2012/0054736	A1	3/2012	Arcese et al.
2012/0072685	A1	3/2012	Otani
2012/0079221	A1	3/2012	Sivasubramanian et al.
2012/0084262	A1	4/2012	Dwarampudi et al.
2012/0084769	A1	4/2012	Adi et al.
2012/0096149	A1	4/2012	Sunkara et al.
2012/0110328	A1	5/2012	Pate et al.
2012/0131295	A1	5/2012	Nakajima
2012/0131578	A1	5/2012	Ciano et al.
2012/0136832	A1	5/2012	Sadhwani
2012/0150815	A1	6/2012	Parfumi
2012/0150818	A1	6/2012	Vijayan Retnamma et al.
2012/0150826	A1	6/2012	Vijayan Retnamma et al.
2012/0151084	A1	6/2012	Stathopoulos et al.
2012/0159232	A1	6/2012	Shimada et al.
2012/0167083	A1	6/2012	Suit
2012/0209812	A1	8/2012	Bezbaruah
2012/0215911	A1	8/2012	Raleigh et al.
2012/0221843	A1	8/2012	Bak et al.
2012/0233285	A1	9/2012	Suzuki
2012/0254119	A1	10/2012	Kumarasamy et al.
2012/0254364	A1	10/2012	Vijayan
2012/0254824	A1	10/2012	Bansold
2012/0278287	A1	11/2012	Wilk
2012/0278571	A1	11/2012	Fleming et al.
2012/0278799	A1	11/2012	Starks et al.
2012/0290802	A1	11/2012	Wade et al.
2012/0324183	A1	12/2012	Chiruvolu et al.
2012/0331248	A1	12/2012	Kono et al.
2013/0019240	A1	1/2013	Tsirkin et al.
2013/0024641	A1	1/2013	Talagala et al.
2013/0024722	A1	1/2013	Kotagiri
2013/0036418	A1	2/2013	Yadappanavar et al.
2013/0042234	A1	2/2013	Deluca et al.
2013/0054533	A1	2/2013	Hao et al.
2013/0061014	A1	3/2013	Prahlad et al.
2013/0074181	A1	3/2013	Singh
2013/0080841	A1	3/2013	Reddy et al.
2013/0086580	A1	4/2013	Simonsen et al.
2013/0111262	A1	5/2013	Taylor et al.
2013/0117744	A1	5/2013	Klein et al.
2013/0167145	A1	6/2013	Krishnamurthy et al.
2013/0173771	A1	7/2013	Ditto et al.
2013/0198828	A1	8/2013	Pendergrass et al.
2013/0204849	A1	8/2013	Chacko
2013/0219069	A1	8/2013	Yellapragada
2013/0227558	A1	8/2013	Du et al.
2013/0232215	A1	9/2013	Gupta et al.
2013/0232480	A1	9/2013	Winterfeldt et al.
2013/0238562	A1	9/2013	Kumarasamy
2013/0262390	A1	10/2013	Kumarasamy et al.
2013/0262638	A1	10/2013	Kumarasamy et al.
2013/0262801	A1	10/2013	Sancheti et al.
2013/0268931	A1	10/2013	O'Hare et al.
2013/0275375	A1	10/2013	Nickolov et al.
2013/0290267	A1	10/2013	Dwarampudi et al.
2013/0311429	A1	11/2013	Agetsuma
2013/0326260	A1	12/2013	Wei et al.
2013/0346709	A1	12/2013	Wang
2014/0006357	A1	1/2014	Davis et al.
2014/0006858	A1	1/2014	Helfman et al.
2014/0007097	A1	1/2014	Chin et al.
2014/0007181	A1	1/2014	Sarin et al.
2014/0052892	A1	2/2014	Klein et al.
2014/0059380	A1	2/2014	Krishnan
2014/0067363	A1	3/2014	Ogren et al.
2014/0075440	A1	3/2014	Prahlad et al.
2014/0089266	A1	3/2014	Une et al.
2014/0095816	A1	4/2014	Hsu et al.
2014/0115285	A1	4/2014	Arcese et al.
2014/0136803	A1	5/2014	Qin
2014/0156684	A1	6/2014	Zaslavsky et al.
2014/0181038	A1	6/2014	Pawar et al.
2014/0181044	A1	6/2014	Pawar et al.
2014/0181046	A1	6/2014	Pawar et al.
2014/0188803	A1	7/2014	James et al.
2014/0196038	A1	7/2014	Kottomtharayil et al.
2014/0196039	A1	7/2014	Kottomtharayil et al.
2014/0201151	A1	7/2014	Kumarasamy et al.
2014/0201157	A1	7/2014	Pawar et al.
2014/0201162	A1	7/2014	Kumarasamy et al.
2014/0201170	A1	7/2014	Vijayan et al.

(56)

References Cited

U.S. PATENT DOCUMENTS

- 2014/0237537 A1 8/2014 Manmohan et al.
 2014/0244610 A1 8/2014 Raman et al.
 2014/0259015 A1 9/2014 Chigusa et al.
 2014/0278530 A1 9/2014 Bruce et al.
 2014/0282514 A1 9/2014 Carson et al.
 2014/0330874 A1 11/2014 Novak et al.
 2014/0337295 A1 11/2014 Haselton et al.
 2014/0344323 A1 11/2014 Pelavin et al.
 2014/0372384 A1 12/2014 Long et al.
 2015/0058382 A1 2/2015 St. Laurent
 2015/0067393 A1 3/2015 Madani et al.
 2015/0074536 A1 3/2015 Varadharajan et al.
 2015/0106557 A1 4/2015 Yu et al.
 2015/0120928 A1 4/2015 Gummaraju et al.
 2015/0121122 A1 4/2015 Towstopiat et al.
 2015/0134607 A1 5/2015 Magdon-Ismail et al.
 2015/0142745 A1 5/2015 Tekade et al.
 2015/0160884 A1 6/2015 Scales et al.
 2015/0161015 A1 6/2015 Kumarasamy et al.
 2015/0163172 A1 6/2015 Mudigonda et al.
 2015/0212895 A1 7/2015 Pawar et al.
 2015/0227438 A1 8/2015 Jaquette
 2015/0227602 A1 8/2015 Ramu
 2015/0242283 A1 8/2015 Simoncelli et al.
 2015/0248333 A1 9/2015 Aravot
 2015/0293817 A1 10/2015 Subramanian et al.
 2015/0317216 A1 11/2015 Hsu et al.
 2015/0324217 A1 11/2015 Shilmover et al.
 2015/0347165 A1 12/2015 Lipchuk et al.
 2015/0347430 A1 12/2015 Ghosh
 2015/0363413 A1 12/2015 Ghosh
 2015/0370652 A1 12/2015 He et al.
 2015/0378758 A1 12/2015 Duggan et al.
 2015/0378771 A1 12/2015 Tarasuk-Levin
 2015/0378833 A1 12/2015 Misra et al.
 2015/0378849 A1 12/2015 Liu et al.
 2015/0381711 A1 12/2015 Singh et al.
 2016/0019224 A1 1/2016 Ahn
 2016/0019317 A1 1/2016 Pawar et al.
 2016/0070623 A1 3/2016 Derk
 2016/0092467 A1 3/2016 Lee et al.
 2016/0127307 A1 5/2016 Jain et al.
 2016/0154709 A1 6/2016 Mitkar et al.
 2016/0170844 A1 6/2016 Long et al.
 2016/0188413 A1 6/2016 Abali et al.
 2016/0202916 A1 7/2016 Cui et al.
 2016/0283335 A1 9/2016 Yao et al.
 2016/0306706 A1 10/2016 Pawar et al.
 2016/0308722 A1 10/2016 Kumarasamy et al.
 2016/0335007 A1 11/2016 Ryu et al.
 2016/0350391 A1 12/2016 Vijayan et al.
 2017/0090972 A1 3/2017 Ryu et al.
 2017/0123939 A1 5/2017 Maheshwari et al.
 2017/0168903 A1 6/2017 Dornemann et al.
 2017/0185488 A1 6/2017 Kumarasamy et al.
 2017/0192866 A1 7/2017 Vijayan et al.
 2017/0193003 A1 7/2017 Vijayan et al.
 2017/0235647 A1 8/2017 Kilaru et al.
 2017/0242871 A1 8/2017 Kilaru et al.
 2017/0249220 A1 8/2017 Kumarasamy et al.
 2017/0262204 A1 9/2017 Dornemann et al.
 2017/0264589 A1 9/2017 Hunt et al.
 2017/0286230 A1 10/2017 Zamir
 2017/0371547 A1 12/2017 Fruchtman et al.
 2018/0089031 A1 3/2018 Dornemann et al.
 2018/0113623 A1 4/2018 Sancheti
 2018/0143880 A1 5/2018 Dornemann
 2018/0181598 A1 6/2018 Pawar et al.
 2018/0246756 A1 8/2018 Abali et al.
 2018/0253192 A1 9/2018 Varadharajan et al.
 2018/0276022 A1 9/2018 Mitkar et al.
 2018/0276083 A1 9/2018 Mitkar et al.
 2018/0276084 A1 9/2018 Mitkar et al.
 2018/0276085 A1 9/2018 Mitkar et al.
 2018/0284986 A1 10/2018 Bhagi
- 2018/0285202 A1 10/2018 Bhagi et al.
 2018/0285353 A1 10/2018 Rao et al.
 2018/0329636 A1 11/2018 Dornemann et al.
 2019/0012339 A1 1/2019 Kumarasamy et al.
 2019/0026187 A1 1/2019 Gulam et al.
 2019/0065069 A1 2/2019 Sancheti et al.
 2019/0090305 A1 3/2019 Hunter et al.
 2019/0340088 A1 11/2019 Sanakkayala et al.
 2019/0347120 A1 11/2019 Kottomtharayil et al.
 2019/0369901 A1 12/2019 Dornemann et al.
 2019/0391742 A1 12/2019 Bhagi et al.
 2020/0034252 A1 1/2020 Mitkar et al.
 2020/0142612 A1 5/2020 Dornemann et al.
 2020/0142782 A1 5/2020 Dornemann
 2020/0142783 A1 5/2020 Dornemann
 2020/0174894 A1 6/2020 Dornemann
 2020/0174895 A1 6/2020 Dornemann
 2020/0183728 A1 6/2020 Deshpande et al.
 2020/0241908 A1 7/2020 Dornemann et al.
 2020/0265024 A1 8/2020 Pawar et al.
 2020/0301891 A1 9/2020 Dornemann
 2020/0327163 A1 10/2020 Pawar et al.
 2020/0334113 A1 10/2020 Sanakkayala et al.
 2020/0334201 A1 10/2020 Pawar et al.
 2021/0089337 A1 3/2021 Kottomtharayil et al.
 2021/0103556 A1 4/2021 Nagrale et al.
 2021/0117294 A1 4/2021 Mitkar et al.
 2021/0240308 A1 8/2021 Varadharajan et al.
 2021/0255937 A1 8/2021 Mitkar et al.
 2022/0083429 A1 3/2022 Pawar et al.
 2022/0206907 A1 6/2022 Bhagi et al.
 2022/0229737 A1 7/2022 Dornemann

FOREIGN PATENT DOCUMENTS

- EP 0467546 A2 1/1992
 EP 0541281 A2 5/1993
 EP 0774715 A1 5/1997
 EP 0809184 A1 11/1997
 EP 0817040 A2 1/1998
 EP 0899662 A1 3/1999
 EP 0981090 A1 2/2000
 WO WO 95/13580 5/1995
 WO WO 99/12098 3/1999
 WO WO 2006/052872 5/2006

OTHER PUBLICATIONS

- Armstead et al., "Implementation of a Campus-wide Distributed Mass Storage Service: The Dream vs. Reality," IEEE, Sep. 11-14, 1995, pp. 190-199.
 Arneson, "Mass Storage Archiving in Network Environments," Digest of Papers, Ninth IEEE Symposium on Mass Storage Systems, Oct. 31, 1988-Nov. 3, 1988, pp. 45-50, Monterey, CA.
 Brandon, J., "Virtualization Shakes Up Backup Strategy," <<http://www.computerworld.com>>, Feb. 21, 2008, 3 pages.
 Cabrera et al., "ADSM: A Multi-Platform, Scalable, Backup and Archive Mass Storage System," Digest of Papers, Compton '95, Proceedings of the 40th IEEE Computer Society International Conference, Mar. 5, 1995-Mar. 9, 1995, pp. 420-427, San Francisco, CA.
 Celesti, et al., "Improving Virtual Machine Migration in Federated Cloud Environments", 2010, pp. 61-67.
 Chan, et al., "An Approach to High Availability for Cloud Servers with Snapshot Mechanism," 2012, pp. 1-6.
 Chen et al., "When Virtual Is Better Than Real", IEEE 2001, pp. 133-138.
 Chervenak, et al., "Protecting File Systems—A Survey of Backup Techniques," 1998, pp. 17-31.
 Chiappetta, Marco, "ESA Enthusiast System Architecture," <<http://hothardware.com/Articles/NVIDIA—ESA—Enthusiast—System—Architecture/>>, Nov. 5, 2007, 2 pages.
 CommVault Systems, Inc., "A CommVault White Paper: VMware Consolidated Backup (VCB) Certification Information Kit," 2007, 23 pages.

(56)

References Cited

OTHER PUBLICATIONS

- CommVault Systems, Inc., "CommVault Solutions—VMware," <<http://www.commvault.com/solutions/vmware/>>, accessed Apr. 30, 2014, 1 page.
- CommVault Systems, Inc., "Enhanced Protection and Manageability of Virtual Servers," Partner Solution Brief, 2008, 6 pages.
- Cully, et al., "Remus: High Availability via Asynchronous Virtual Machine Replication", 2008, pp. 161-174.
- Data Protection for Large VMware and Vblock Environments Using EMC Avamar Applied Technology, Nov. 2010, EMC Corporation, 26 pages.
- Davis, D., "3 VMware Consolidated Backup (VCB) Utilities You Should Know," Petri IT Knowledgebase, <<http://www.petri.co.il/vmware-consolidated-backup-utilities.htm>>, Nov. 16, 2007, 3 pages.
- Davis, D., "Understanding VMware VMX Configuration Files," Petri IT Knowledgebase, <http://www.petri.co.il/virtual_vmware_vmx_configuration_files.htm>, Nov. 16, 2007, 3 pages.
- Davis, D., "VMware Server & Workstation Disk Files Explained," Petri IT Knowledgebase, <http://www.petri.co.il/virtual_vmware_files_explained.htm>, May 3, 2008, 3 pages.
- Davis, D., "VMware Versions Compared," Petri IT Knowledgebase, <http://www.petri.co.il/virtual_vmware_versions_compared.htm>, Nov. 16, 2007, 3 pages.
- Deng, et al., "Fast Saving and Restoring Virtual Machines with Page Compression", 2011, pp. 150-157.
- Eitel, "Backup and Storage Management in Distributed Heterogeneous Environments," IEEE, Jun. 12-16, 1994, pp. 124-126.
- Eldos Callback File System product information from <https://www.eldos.com/clients/104-345.php> retrieved on Dec. 30, 2016 in 2 pages.
- Eldos Usermode filesystem for your Windows applications—Callback File System® (CBFS®)—Create and manage virtual filesystems and disks from your Windows applications retrieved from <https://eldos.com/cbfs> on Dec. 30, 2016 in 4 pages.
- Fraser, et al., "Safe Hardware Access With the Xen Virtual Machine Monitor", 1st Workshop on Operating System and Architectural Support for the demand IT Infrastructure (OASIS), 2004, pp. 1-10.
- Gait, "The Optical File Cabinet: A Random-Access File System for Write-Once Optical Disks," IEEE Computer, vol. 21, No. 6, pp. 11-22 (Jun. 1988) (see in particular figure 5 in p. 15 and recitation in claim 5).
- Galan et al. "Service Specification in Cloud Environments Based on Extension to Oper Standards" COMSWARE Jun. 16-19, 2009 Dublin, Ireland ACM.
- Gibson, et al., "Implementing Preinstallation Environment Media for Use in User Support," 2007, pp. 129-130.
- Granger, et al., "Survivable Storage Systems", 2001, pp. 184-195.
- Gupta, et al., "GPFS-SNC: An enterprise storage framework for virtual-machine clouds", 2011, pp. 1-10.
- Haselhorst, et al., "Efficient Storage Synchronization for Live Migration in Cloud Infrastructures", 2011, pp. 511-518.
- Hirofuchio, Takahiro et al., "A live storage migration mechanism over wan and its performance evaluation," 2009, pp. 67-74.
- Hirofuchi, et al., "Enabling Instantaneous Relocation of Virtual Machines with a Lightweight VMM Extension", 2010, pp. 73-83.
- Hu, et al., "Virtual Machine based Hot-spare Fault-tolerant System", 2009, pp. 429-432.
- Hu, Wenjin et al., "A Quantitative Study of Virtual Machine Live Migration," 2013, pp. 1-10.
- Huff, "Data Set Usage Sequence Number," IBM Technical Disclosure Bulletin, vol. 24, No. 5, Oct. 1981 New York, US, pp. 2404-2406.
- Ibrahim, Shadi et al., "CLOUDLET: Towards MapReduce Implementation on Virtual Machines," 2009, pp. 65-66.
- Ismail et al., Architecture of Scalable Backup Service For Private Cloud, IEEE 2013, pp. 174-179.
- Jander, "Launching Storage-Area Net," Data Communications, US, McGraw Hill, NY, vol. 27, No. 4 (Mar. 21, 1998), pp. 64-72.
- Javaraiah, et al., "Backup for Cloud and Disaster Recovery for Consumers and SMBs," 2008, pp. 1-3.
- Jhavar et al., "Fault Tolerance Management in Cloud Computing: A System-Level Perspective", IEEE Systems Journal 7.2, 2013, pp. 288-297.
- Jo, et al., "Efficient Live Migration of Virtual Machines Using Shared Storage", 2013, pp. 1-10.
- Kashyap "RLC—A Reliable approach to Fast and Efficient Live Migration of Virtual Machines in the Clouds" IEEE 2014 IEEE Computer Society.
- Kim, et al., "Availability Modeling and Analysis of a Virtualized System," 2009, pp. 365-371.
- Kuo, et al., "A Hybrid Cloud Storage Architecture for Service Operational High Availability", 2013, pp. 487-492.
- Li et al. "Comparing Containers versus Virtual Machines for Achieving High Availability" 2015 IEEE.
- Liang, et al., "A virtual disk environment for providing file system recovery", 2006, pp. 589-599.
- Lu et al. "Virtual Machine Memory Access Tracing with Hypervisor Exclusive Cache", Usenix Annual Technical Conference, 2007, pp. 29-43.
- Mao, et al., "Read-Performance Optimization for Deduplication-Based Storage Systems in the Cloud", 2014, pp. 1-22.
- Microsoft Corporation, "How NTFS Works," Windows Server TechCenter, <<http://technet2.microsoft.com/windowsserver/en/library/8cc5891d-bf8e-4164-862d-dac5418c5948> . . . >, updated Mar. 28, 2003, internet accessed Mar. 26, 2008, 26 pages.
- Migrate a Virtual Machine with Storage vMotion in the vSphere Client. <http://pubs.vmware.com/vsphere-51/advanced/print/jsp?topic=/com.vmware.vsphere.vcent> . . . Retrieved Aug. 12, 2014; 2 pages.
- Nance et al., "Virtual Machine Introspection: Observation or Interference?", 2008 IEEE.
- Ng, Chun-Ho et al., "Live Deduplication Storage of Virtual Machine Images in an Open-Source Cloud," 2011, pp. 80-99.
- Nicolae, Bogdan et al., "A Hybrid Local Storage Transfer Scheme for Live Migration of I/O Intensive Workloads," 2012, pp. 85-96.
- Reingold, B. et al., "Cloud Computing: The Intersection of Massive Scalability, Data Security and Privacy (Part I)," LegalWorks, a Thomson Business, Jun. 2009, 5 pages.
- Reingold, B. et al., "Cloud Computing: Industry and Government Developments (Part II)," LegalWorks, Sep. 2009, 5 pages.
- Reingold, B. et al., "Cloud Computing: Whose Law Governs the Cloud? (Part III)," LegalWorks, Jan.-Feb. 2010, 6 pages.
- Rosenblum et al., "The Design and Implementation of a Log-Structured File System," Operating Systems Review SIGOPS, vol. 25, No. 5, New York, US, pp. 1-15 (May 1991).
- Sanbarrow.com, "Disktype-table," <<http://sanbarrow.com/vmdk/disktypes.html>>, internet accessed on Apr. 30, 2014, 4 pages.
- Sanbarrow.com, "Files Used by a VM," <<http://sanbarrow.com/vmx/vmx-files-used-by-a-vm.html>>, internet accessed on Apr. 30, 2014, 1 page.
- Sanbarrow.com, "Monolithic Versus Split Disks," <<http://sanbarrow.com/vmdk/monolithicversussplit.html>>, internet accessed on Jul. 14, 2008, 2 pages.
- Somasundaram et al., Information Storage and Management. 2009, pp. 251-281.
- Tran, et al., "Efficient Cooperative Backup with Decentralized Trust Management", 2012, pp. 1-25.
- Travostino, et al., "Seamless live migration of virtual machines over the MAN/WAN", 2006, pp. 901-907.
- Tudoran, Radu et al., "Adaptive File Management for Scientific Workflows on the Azure Cloud," 2013, pp. 273-281.
- Vaghani, "Virtual Machine File System", 2010, pp. 57-70.
- VMware, Inc., "VMware Solution Exchange (VSX)" <<http://www.vmware.com/appliances/learn/ovf.html>>, 2014, 3 pages.
- VMware, Inc., "OVF, Open Virtual Machine Format Specification, version 0.9," White Paper, <<http://www.vmware.com>>, Sep. 7, 2007, 50 pages.
- VMware, Inc., "The Open Virtual Machine Format Whitepaper for OVF Specification, version 0.9," White Paper, <<http://www.vmware.com>>, 2007, 16 pages.
- VMware, Inc., "Understanding VMware Consolidated Backup," White Paper, <<http://www.vmware.com>>, accessed Apr. 30, 2014, 11 pages.

(56)

References Cited

OTHER PUBLICATIONS

- VMware, Inc., "Using VMware Infrastructure for Backup and Restore," Best Practices, <<http://www.vmware.com>>, accessed Apr. 30, 2014, 20 pages.
- VMware, Inc., "Virtual Disk API Programming Guide," <<http://www.vmware.com>>, Revision 20080411, 2008, 44 pages.
- VMware, Inc., "Virtual Disk Format 1.1," VMware Technical Note, <<http://www.vmware.com>>, Revision 20071113, Version 1.1, 2007, 18 pages.
- VMware, Inc., "Virtualized iSCSI SANS: Flexible, Scalable Enterprise Storage for Virtual Infrastructures," White Paper, <<http://www.vmware.com>>, Mar. 2008, 13 pages.
- VMware, Inc., "Virtual Machine Backup Guide, ESX Server 3.5, ESX Server 3i version 3.5, VirtualCenter 2.5," <<http://www.vmware.com>>, updated Feb. 21, 2008, 78 pages.
- VMware, Inc., "Virtual Machine Backup Guide, ESX Server 3.0.1 and VirtualCenter 2.0.1," <<http://www.vmware.com>>, updated Nov. 21, 2007, 74 pages.
- VMware, Inc., "VMware Consolidated Backup," Product Datasheet, <<http://www.vmware.com>>, 2009, 2 pages.
- VMware, Inc., "VMware Consolidated Backup, Improvements in Version 3.5," Information Guide, <<http://www.vmware.com>>, accessed Apr. 30, 2014, 11 pages.
- VMware, Inc., "VMware ESX 3.5," Product Datasheet, <<http://www.vmware.com>>, 2008, 4 pages.
- VMware, Inc., "VMware GSX Server 3.2, Disk Types: Virtual and Physical," <http://www.vmware.com/support/gsx3/doc/disks_types_gsx.html>, 2008, 2 pages.
- VMware, Inc., "VMware OVF Tool," Technical Note, <<http://www.vmware.com>>, 2007, 4 pages.
- VMware, Inc., "VMware Workstation 5.0, Snapshots in a Process Tree," <http://www.vmware.com/support/ws5/doc/ws_preserve_sshot_tree.html>, accessed Apr. 30, 2014, 1 page.
- VMware, Inc., "VMware Workstation 5.0, Snapshots in a Linear Process," <http://www.vmware.com/support/ws5/doc/ws_preserve_sshot_linear.html>, internet accessed on 2014, 1 page.
- VMware, Inc., "VMware Workstation 5.5, What Files Make Up a Virtual Machine?" <http://www.vmware.com/support/ws55/doc/ws_learning_files_in_a_vm.html>, 2014, 2 pages.
- VMware Storage VMotion—Non-Disruptive Live Migration for Virtual Machine Storage Disk Files. Copyright 2009 VMware, Inc.; 2 pages.
- Vrable, et al., "Cumulus: Filesystem Backup to the Cloud", 2009, pp. 1-28.
- VSphere Storage vMotion: Storage Management & Virtual Machine Migration. <http://www.vmware.com/products/vsphere/features/storage-vmotion> Retrieved Aug. 12, 2014; 6 pages.
- Wikipedia, "Cloud computing," <<http://en.wikipedia.org/wiki/Cloud-computing>>, 2009, 11 pages.
- Wikipedia, "Cluster (file system)," <http://en.wikipedia.org/wiki/Cluster_%28file_system%29>, Sep. 2, 2008, 1 page.
- Wikipedia, "Cylinder-head-sector," <<http://en.wikipedia.org/wiki/Cylinder-head-sector>>, Jan. 4, 2009, 6 pages.
- Wikipedia, "File Allocation Table," <http://en.wikipedia.org/wiki/File_Allocation_Table>, Dec. 3, 2008, 12 pages.
- Wikipedia, "Logical Disk Manager," <http://en.wikipedia.org/wiki/Logical_Disk_Manager>, Nov. 16, 2007, 3 pages.
- Wikipedia, "Logical Volume Management," <http://en.wikipedia.org/wiki/Logical_volume_management>, Oct. 27, 2008, 3 pages.
- Wikipedia, "Storage Area Network," <http://en.wikipedia.org/wiki/Storage_area_network>, Dec. 5, 2008, 5 pages.
- Wikipedia, "Virtualization," <<http://en.wikipedia.org/wiki/Virtualization>>, Apr. 29, 2014, 7 pages.
- Wood, et al., "Disaster Recovery as a Cloud Service: Economic Benefits & Deployment Challenges", 2010, pp. 1-7.
- Yang, et al., "Toward Reliable Data Delivery for Highly Dynamic Mobile Ad Hoc Networks," 2012, pp. 111-124.
- Yang, et al., "TRAP-Array: A Disk Array Architecture Providing Timely Recovery to Any Point-in-time," 2006, pp. 1-12.
- Yoshida et al., "Orthros: A High-Reliability Operating System with Transmigration of Processes," 2013, pp. 318-327.
- Zhao, et al., "Adaptive Distributed Load Balancing Algorithm based on Live Migration of Virtual Machines in Cloud", 2009, pp. 170-175.
- Zhao, et al., Supporting Application—Tailored Grid File System Sessions with WSRF-Based Services, Advanced Computing and Information Systems Laboratory (ACIS), pp. 24-33.
- International Search Report and Written Opinion for PCT/US2011/054374, dated May 2, 2012, 7 pages.
- International Preliminary Report on Patentability and Written Opinion for PCT/US2011/054374, dated Apr. 2, 2013, 9 pages.
- Kottomtharayil, et al., U.S. Appl. No. 14/148,465, Published as 2014/0196038 A1 Abandoned, filed Jan. 6, 2014, Virtual Machine Management in a Data Storage System.
- Kottomtharayil, et al., U.S. Appl. No. 14/148,549, Published as 2014/0196039 A1 Abandoned, filed Jan. 6, 2014, Virtual Machine Categorization System and Method.
- Kottomtharayil, et al., U.S. Appl. No. 14/148,507 Published as 2014/0196056 A1 Now U.S. Pat. No. 9,703,584, filed Jan. 6, 2014, Virtual Server Agent Load Balancing.
- Kottomtharayil, et al., U.S. Appl. No. 15/196,739 Published as 2016/0306642 A1 Now U.S. Pat. No. 9,977,687, filed Jun. 29, 2016, Virtual Server Agent Load Balancing.
- Kottomtharayil, et al., U.S. Appl. No. 15/960,401 Published as 2018/0307510 A1 Now U.S. Pat. No. 10,474,483, filed Apr. 23, 2018, Virtual Server Agent Load Balancing.
- Kottomtharayil, et al., U.S. Appl. No. 16/518,107 Published as 2019/0347120 A1, filed Jul. 22, 2019, Virtual Machine Load Balancing.




FIG. 1A




FIG. 1B




FIG. 1C




FIG. 1D




FIG. 1E




FIG. 1H




FIG. 2

VIRTUAL MACHINE JOB ALLOCATION




FIG. 3

VIRTUAL MACHINE GROUPING




FIG. 4

VIRTUAL SERVER AGENT LOAD BALANCING
(BACKUP)

500
↘




FIG. 5

VIRTUAL MACHINE LOAD BALANCING

RELATED APPLICATIONS

This disclosure is a continuation of U.S. patent application Ser. No. 16/518,107, filed on Jul. 22, 2019, and titled “VIRTUAL MACHINE LOAD BALANCING,” which is a continuation of U.S. patent application Ser. No. 15/960,401, filed on Apr. 23, 2018, and titled “VIRTUAL SERVER AGENT LOAD BALANCING,” which is a continuation of U.S. patent application Ser. No. 15/196,739, filed on Jun. 29, 2016, and titled “VIRTUAL SERVER AGENT LOAD BALANCING,” the disclosure of which is a continuation of U.S. patent application Ser. No. 14/148,507, filed on Jan. 6, 2014, and titled “VIRTUAL SERVER AGENT LOAD BALANCING,” which claims the benefit of priority under 35 U.S.C. § 119(e) of U.S. Provisional Patent Application No. 61/750,255, filed on Jan. 8, 2013, and titled “VIRTUAL MACHINE MANAGEMENT IN A DATA STORAGE SYSTEM,” the disclosures of which are hereby incorporated by reference in their entireties. Further, this disclosure is related to the following disclosures that were filed on Jan. 6, 2014, the same date as the parent disclosure, and which are hereby incorporated by reference in their entirety herein: U.S. application Ser. No. 14/148,465, titled “VIRTUAL MACHINE MANAGEMENT IN A DATA STORAGE SYSTEM” and U.S. application Ser. No. 14/148,549, titled “VIRTUAL MACHINE CATEGORIZATION SYSTEM AND METHOD.” This disclosure is also related to another disclosure filed on the same day as this disclosure: U.S. patent application Ser. No. 17/118,344, and titled “VIRTUAL SERVER AGENT LOAD BALANCING”.

BACKGROUND

Businesses worldwide recognize the commercial value of their data and seek reliable, cost-effective ways to protect the information stored on their computer networks while minimizing impact on productivity. Protecting information is often part of a routine process that is performed within an organization.

A company might back up critical computing systems such as databases, file servers, web servers, and so on as part of a daily, weekly, or monthly maintenance schedule. The company may similarly protect computing systems used by each of its employees, such as those used by an accounting department, marketing department, engineering department, and so forth.

Given the rapidly expanding volume of data under management, companies also continue to seek innovative techniques for managing data growth, in addition to protecting data. For instance, companies often implement migration techniques for moving data to lower cost storage over time and data reduction techniques for reducing redundant data, pruning lower priority data, etc.

Enterprises also increasingly view their stored data as a valuable asset. Along these lines, customers are looking for solutions that not only protect and manage, but also leverage their data. For instance, solutions providing data analysis capabilities, improved data presentation and access features, and the like, are in increasing demand.

In certain environments, data storage operations can be implemented with the use of virtual machines. As such virtual machines are generally allocated physical resources for operational purposes, excessive utilization numbers of virtual machines can limit a system’s available resources. Furthermore, inefficient allocation/use of virtual machines

can affect a system’s operational capacity. Therefore, effective virtual machine management may be a concern in data storage systems.

SUMMARY

For purposes of summarizing the disclosure, certain aspects, advantages and novel features of the inventions have been described herein. It is to be understood that not necessarily all such advantages may be achieved in accordance with any particular embodiment of the inventions disclosed herein. Thus, the inventions disclosed herein may be embodied or carried out in a manner that achieves or optimizes one advantage or group of advantages as taught herein without necessarily achieving other advantages as may be taught or suggested herein.

In large enterprise environments, numerous virtual machines (VMs) may be instantiated. Often, a number of virtual machines may remain unused or may be left running after a user or system has completed accessing the virtual machine. This proliferation of virtual machines can result in wasted resources. Further, managing backup of a primary storage system to a secondary storage system can be complicated and require significant resources due at least in part to the proliferation of virtual machines.

In order to address the above and other challenges relating to virtual machine proliferation, an information management system is provided that includes a number of virtual server agents (VSAs) and a VSA coordinator that can reduce VM proliferation. A job or task to be performed by a virtual machine may instead be provided to a VSA, which can identify an available VM to perform the task. If a VM is not available, a new VM may be created within a set of VM host systems. However, if an existing VM has resources available to process the task, the VSA may provide the task to the existing VM. Advantageously, in certain embodiments, by providing tasks to existing VMs, the number of VMs instantiated in the information management system are reduced.

In some implementations, VMs or VM host systems are grouped based on one or more factors. For example, VM host systems may be grouped based on the capabilities available to the VM host system and/or VMs hosted by the VM host system. The groups of VMs or VM host systems may then be assigned or associated with one or more VSAs to help manage job allocations to VMs within the group.

In some embodiments, the VSA coordinator facilitates load balancing of the VSAs. During a backup process, the VSA coordinator may identify a number of VSAs available to help backup a set of VMs from a primary storage system to a secondary storage system. The VSA coordinator may allocate the VMs among the available VSAs based on a number of communication streams between each VSA and systems of the secondary storage system. Further, in some cases, the VSA coordinator determines an allocation of the VMs based on characteristics of the VMs to be backed up, such as the type of VM or the size of the VM.

Certain embodiments described herein include a method of reducing virtual machine proliferation. The method may include receiving a job request at a virtual server agent. This virtual server agent may include computer hardware. Further, the method may include determining a load for each virtual machine from a set of virtual machines. The set of virtual machines may be at least partially managed by the virtual server agent. In addition, the method may include determining whether the load of at least one virtual machine from the set of virtual machines is below a threshold load. In response to determining that the load of at least one

virtual machine from the set of virtual machines is below the threshold load, the method may include selecting a virtual machine from a set of virtual machines with a load that is below the threshold load and assigning a job associated with the job request to the selected virtual machine. Further, in response to determining that no virtual machine from the set of virtual machines is below a threshold load, the method can include initiating creation of a new virtual machine and assigning the job associated with the job request to the new virtual machine.

In some embodiments, a system for reducing virtual machine proliferation is disclosed. The system can include a virtual server agent comprising computer hardware. The virtual server agent may be configured to receive a job request and to access load information for each virtual machine from a set of virtual machines assigned to the virtual server agent. Further, the virtual server agent may identify, based at least partially on the load information for each virtual machine, a subset of virtual machines from the set of virtual machines with a load below a threshold load. In addition, the virtual server agent can select a virtual machine from the subset of virtual machines and assign a job associated with the job request to the selected virtual machine when the subset of virtual machines is a non-empty set.

Certain embodiments described herein include a method of grouping virtual machines. The method may be performed by a virtual server agent coordinator comprising computer hardware. The method can include identifying a set of virtual machine provider systems in a primary storage system. Each of the virtual machine provider systems may include a virtual machine monitor and may be configured to host a set of virtual machines. Further, the method can include accessing metadata for each of the virtual machine provider systems from the set of virtual machine provider systems. In addition, the method may include grouping the virtual machine provider systems into one or more groups based at least partially on the metadata for each of the virtual machine provider systems. Moreover, the method can include assigning a set of virtual server agents to each group of virtual machine provider systems. Each virtual server agent may be configured to backup data from at least one virtual machine in the primary storage system to a secondary storage system.

In some embodiments, a system for grouping virtual machines is disclosed. The system may include a virtual server agent coordinator comprising computer hardware. The virtual server agent coordinator may be configured to identify a set of virtual machine provider systems in a primary storage system. Each of the virtual machine provider systems may include a virtual machine monitor and may be configured to host a set of virtual machines. Further, the virtual server agent coordinator may access metadata for each of the virtual machine provider systems from the set of virtual machine provider systems. In addition, the virtual server agent coordinator can group the virtual machine provider systems into one or more groups based at least partially on the metadata for each of the virtual machine provider systems. Moreover, the virtual server agent coordinator can assign a set of virtual server agents to each group of virtual machine provider systems. Each virtual server agent may be configured to backup data from at least one virtual machine in the primary storage system to a secondary storage system.

Certain embodiments described herein include a method of virtual server agent load balancing. The method may be performed by a virtual server agent coordinator comprising

computer hardware. The method may include identifying a set of virtual machines for backup to a secondary storage system. The set of virtual machines may be hosted by a set of virtual machine provider systems. Further, the set of virtual machine provider systems may be included in a primary storage system. The method may also include identifying a set of virtual server agents available to backup data from the set of virtual machines to the secondary storage system. Further, the method may include determining a number of data streams available to each virtual server agent from the set of virtual server agents. Moreover, the method may include distributing the set of virtual machines among the set of virtual server agents based at least partially on the number of data streams available to each of the virtual server agents.

In some embodiments, a system for virtual server agent load balancing is disclosed. The system may include a virtual server agent coordinator comprising computer hardware. Further, the virtual server agent coordinator may be configured to identify a set of virtual machines for backup to a secondary storage system. The set of virtual machines may be hosted by a set of virtual machine provider systems. Further, the set of virtual machine provider systems may be included in a primary storage system. Moreover, the virtual server agent coordinator may identify a set of virtual server agents available to backup data from the set of virtual machines to the secondary storage system. In addition, the virtual server agent coordinator may determine a number of data streams available to each virtual server agent from the set of virtual server agents. Moreover, the virtual server agent coordinator can distribute the set of virtual machines among the set of virtual server agents based at least partially on the number of data streams available to each of the virtual server agents.

BRIEF DESCRIPTION OF THE DRAWINGS

Throughout the drawings, reference numbers are re-used to indicate correspondence between referenced elements. The drawings are provided to illustrate embodiments of the inventive subject matter described herein and not to limit the scope thereof.

FIG. 1A is a block diagram illustrating an exemplary information management system.

FIG. 1B is a detailed view of a primary storage device, a secondary storage device, and some examples of primary data and secondary copy data.

FIG. 1C is a block diagram of an exemplary information management system including a storage manager, one or more data agents, and one or more media agents.

FIG. 1D is a block diagram illustrating a scalable information management system.

FIG. 1E illustrates certain secondary copy operations according to an exemplary storage policy.

FIGS. 1F-1H are block diagrams illustrating suitable data structures that may be employed by the information management system.

FIG. 2 is a block diagram illustrating an example of a scalable information management system.

FIG. 3 illustrates a flowchart for an example virtual machine job allocation process.

FIG. 4 illustrates a flowchart for an example virtual machine grouping process.

FIG. 5 illustrates a flowchart for an example virtual server agent load balancing process.

DETAILED DESCRIPTION

Systems and methods are described herein for reducing virtual machine (VM) proliferation, grouping virtual

machines, and load balancing virtual server agents (VSAs). Examples of such systems and methods are discussed in further detail herein, e.g., with respect to FIGS. 2-5. Moreover, it will be appreciated that reducing VM proliferation, grouping VMs, and load balancing VSAs may be implemented by information management systems, such as those that will now be described with respect to FIGS. 1A-1H. Further, as will be described, the components for implementing a reduction in VM proliferation, grouping VMs, and load balancing VSAs can be incorporated into such systems.

System Overview

The system and methods described with respect to FIGS. 1A-1H can be used for implementing a reduction in virtual machine proliferation. In some embodiments, one or more VSAs assigned to a group of virtual machine hosts (e.g., client computing devices and server computing devices) may determine the availability of a VM to perform a task or job. Thus, tasks may be assigned to existing VMs instead of creating a new VM to perform the task. Furthermore, systems and methods for grouping virtual machine hosts are also disclosed herein. In some embodiments, a VSA coordinator may determine a grouping of VMs or VM hosts based on one or more factors associated with the VMs or the VM hosts, such as VM type or geographical location of the VM hosts. The VSA coordinator may also assign one or more VSAs to facilitate managing the group of VM hosts. In some embodiments, the VSA coordinators may facilitate load balancing of VSAs during operation, such as during a backup operation, a restore operation, or any other operation between a primary storage system and a secondary storage system.

Information Management System Overview

With the increasing importance of protecting and leveraging data, organizations simply cannot afford to take the risk of losing critical data. Moreover, runaway data growth and other modern realities make protecting and managing data an increasingly difficult task. There is therefore a need for efficient, powerful, and user-friendly solutions for protecting and managing data.

Depending on the size of the organization, there are typically many data production sources which are under the purview of tens, hundreds, or even thousands of employees or other individuals. In the past, individual employees were sometimes responsible for managing and protecting their data. A patchwork of hardware and software point solutions has been applied in other cases. These solutions were often provided by different vendors and had limited or no interoperability.

Certain embodiments described herein provide systems and methods capable of addressing these and other shortcomings of prior approaches by implementing unified, organization-wide information management. FIG. 1A shows one such information management system **100**, which generally includes combinations of hardware and software configured to protect and manage data and metadata generated and used by the various computing devices in the information management system **100**.

The organization which employs the information management system **100** may be a corporation or other business entity, non-profit organization, educational institution, household, governmental agency, or the like.

Generally, the systems and associated components described herein may be compatible with and/or provide some or all of the functionality of the systems and corresponding components described in one or more of the following U.S. patents and patent application publications

assigned to CommVault Systems, Inc., each of which is hereby incorporated in its entirety by reference herein:

U.S. Pat. No. 8,285,681, entitled "Data Object Store and Server for a Cloud Storage Environment, Including Data Deduplication and Data Management Across Multiple Cloud Storage Sites";

U.S. Pat. No. 8,307,177, entitled "Systems And Methods For Management Of Virtualization Data";

U.S. Pat. No. 7,035,880, entitled "Modular Backup and Retrieval System Used in Conjunction With a Storage Area Network";

U.S. Pat. No. 7,343,453, entitled "Hierarchical Systems and Methods for Providing a Unified View of Storage Information";

U.S. Pat. No. 7,395,282, entitled "Hierarchical Backup and Retrieval System";

U.S. Pat. No. 7,246,207, entitled "System and Method for Dynamically Performing Storage Operations in a Computer Network";

U.S. Pat. No. 7,747,579, entitled "Metabase for Facilitating Data Classification";

U.S. Pat. No. 8,229,954, entitled "Managing Copies of Data";

U.S. Pat. No. 7,617,262, entitled "System and Methods for Monitoring Application Data in a Data Replication System";

U.S. Pat. No. 7,529,782, entitled "System and Methods for Performing a Snapshot and for Restoring Data";

U.S. Pat. No. 8,230,195, entitled "System And Method For Performing Auxiliary Storage Operations";

U.S. Pat. No. 7,315,923, entitled "System And Method For Combining Data Streams In Pipelined Storage Operations In A Storage Network";

U.S. Pat. No. 8,364,652, entitled "Content-Aligned, Block-Based Deduplication";

U.S. Pat. Pub. No. 2006/0224846, entitled "System and Method to Support Single Instance Storage Operations";

U.S. Pat. No. 8,578,120, entitled "Block-Level Single Instancing";

U.S. Pat. Pub. No. 2009/0319534, entitled "Application-Aware and Remote Single Instance Data Management";

U.S. Pat. Pub. No. 2012/0150826, entitled "Distributed Deduplicated Storage System";

U.S. Pat. Pub. No. 2012/0150818, entitled "Client-Side Repository in a Networked Deduplicated Storage System";

U.S. Pat. No. 8,170,995, entitled "Method and System for Offline Indexing of Content and Classifying Stored Data";

U.S. Pat. No. 7,107,298, entitled "System And Method For Archiving Objects In An Information Store";

U.S. Pat. No. 8,230,195, entitled "System And Method For Performing Auxiliary Storage Operations";

U.S. Pat. No. 8,229,954, entitled "Managing Copies Of Data"; and

U.S. Pat. No. 8,156,086, entitled "Systems And Methods For Stored Data Verification".

The information management system **100** can include a variety of different computing devices. For instance, as will be described in greater detail herein, the information management system **100** can include one or more client computing devices **102** and secondary storage computing devices **106**.

Computing devices can include, without limitation, one or more: workstations, personal computers, desktop comput-

ers, or other types of generally fixed computing systems such as mainframe computers and minicomputers.

Other computing devices can include mobile or portable computing devices, such as one or more laptops, tablet computers, personal data assistants, mobile phones (such as smartphones), and other mobile or portable computing devices such as embedded computers, set top boxes, vehicle-mounted devices, wearable computers, etc. Computing devices can include servers, such as mail servers, file servers, database servers, and web servers.

In some cases, a computing device includes virtualized and/or cloud computing resources. For instance, one or more virtual machines may be provided to the organization by a third-party cloud service vendor. Or, in some embodiments, computing devices can include one or more virtual machine(s) running on a physical host computing device (or “host machine”) operated by the organization. As one example, the organization may use one virtual machine as a database server and another virtual machine as a mail server, both virtual machines operating on the same host machine.

A virtual machine includes an operating system and associated virtual resources, and is hosted simultaneously with another operating system on a physical host computer (or host machine). A hypervisor (typically software, and also known in the art as a virtual machine monitor or a virtual machine manager or “VMM”) sits between the virtual machine and the hardware of the physical host computer. One example of hypervisor as virtualization software is ESX Server, by VMware, Inc. of Palo Alto, Calif.; other examples include Microsoft Virtual Server and Microsoft Windows Server Hyper-V, both by Microsoft Corporation of Redmond, Wash., and Sun xVM by Oracle America Inc. of Santa Clara, Calif. In some embodiments, the hypervisor may be firmware or hardware or a combination of software and/or firmware and/or hardware.

The hypervisor provides to each virtual operating system virtual resources, such as a virtual processor, virtual memory, a virtual network device, and a virtual disk. Each virtual machine has one or more virtual disks. The hypervisor typically stores the data of virtual disks in files on the file system of the physical host computer, called virtual machine disk files (in the case of VMware virtual servers) or virtual hard disk image files (in the case of Microsoft virtual servers). For example, VMware’s ESX Server provides the Virtual Machine File System (VMFS) for the storage of virtual machine disk files. A virtual machine reads data from and writes data to its virtual disk much the same way that an actual physical machine reads data from and writes data to an actual disk.

Examples of techniques for implementing information management techniques in a cloud computing environment are described in U.S. Pat. No. 8,285,681, which is incorporated by reference herein. Examples of techniques for implementing information management techniques in a virtualized computing environment are described in U.S. Pat. No. 8,307,177, also incorporated by reference herein.

The information management system **100** can also include a variety of storage devices, including primary storage devices **104** and secondary storage devices **108**, for example. Storage devices can generally be of any suitable type including, without limitation, disk drives, hard-disk arrays, semiconductor memory (e.g., solid state storage devices), network attached storage (NAS) devices, tape libraries or other magnetic, non-tape storage devices, optical media storage devices, DNA/RNA-based memory technology, combinations of the same, and the like. In some embodiments, storage devices can form part of a distributed

file system. In some cases, storage devices are provided in a cloud (e.g., a private cloud or one operated by a third-party vendor). A storage device in some cases comprises a disk array or portion thereof.

The illustrated information management system **100** includes one or more client computing device **102** having at least one application **110** executing thereon, and one or more primary storage devices **104** storing primary data **112**. The client computing device(s) **102** and the primary storage devices **104** may generally be referred to in some cases as a primary storage subsystem **117**. A computing device in an information management system **100** that has a data agent **142** installed on it is generally referred to as a client computing device **102** (or, in the context of a component of the information management system **100** simply as a “client”).

Depending on the context, the term “information management system” can refer to generally all of the illustrated hardware and software components. Or, in other instances, the term may refer to only a subset of the illustrated components.

For instance, in some cases, the information management system **100** generally refers to a combination of specialized components used to protect, move, manage, manipulate, analyze, and/or process data and metadata generated by the client computing devices **102**. However, the information management system **100** in some cases does not include the underlying components that generate and/or store the primary data **112**, such as the client computing devices **102** themselves, the applications **110** and operating system residing on the client computing devices **102**, and the primary storage devices **104**. As an example, “information management system” may sometimes refer to one or more of the following components and corresponding data structures: storage managers, data agents, and media agents. These components will be described in further detail below.

Client Computing Devices

There are typically a variety of sources in an organization that produce data to be protected and managed. As just one illustrative example, in a corporate environment such data sources can be employee workstations and company servers such as a mail server, a web server, or the like. In the information management system **100**, the data generation sources include the one or more client computing devices **102**.

The client computing devices **102** may include any of the types of computing devices described above, without limitation, and in some cases the client computing devices **102** are associated with one or more users and/or corresponding user accounts, of employees or other individuals.

The information management system **100** generally addresses and handles the data management and protection needs for the data generated by the client computing devices **102**. However, the use of this term does not imply that the client computing devices **102** cannot be “servers” in other respects. For instance, a particular client computing device **102** may act as a server with respect to other devices, such as other client computing devices **102**. As just a few examples, the client computing devices **102** can include mail servers, file servers, database servers, and web servers.

Each client computing device **102** may have one or more applications **110** (e.g., software applications) executing thereon which generate and manipulate the data that is to be protected from loss and managed.

The applications **110** generally facilitate the operations of an organization (or multiple affiliated organizations), and can include, without limitation, mail server applications

(e.g., Microsoft Exchange Server), file server applications, mail client applications (e.g., Microsoft Exchange Client), database applications (e.g., SQL, Oracle, SAP, Lotus Notes Database), word processing applications (e.g., Microsoft Word), spreadsheet applications, financial applications, presentation applications, browser applications, mobile applications, entertainment applications, and so on.

The client computing devices **102** can have at least one operating system (e.g., Microsoft Windows, Mac OS X, iOS, IBM z/OS, Linux, other Unix-based operating systems, etc.) installed thereon, which may support or host one or more file systems and other applications **110**.

As shown, the client computing devices **102** and other components in the information management system **100** can be connected to one another via one or more communication pathways **114**. The communication pathways **114** can include one or more networks or other connection types including as any of following, without limitation: the Internet, a wide area network (WAN), a local area network (LAN), a Storage Area Network (SAN), a Fibre Channel connection, a Small Computer System Interface (SCSI) connection, a virtual private network (VPN), a token ring or TCP/IP based network, an intranet network, a point-to-point link, a cellular network, a wireless data transmission system, a two-way cable system, an interactive kiosk network, a satellite network, a broadband network, a baseband network, a neural network, a mesh network, an ad hoc network, other appropriate wired, wireless, or partially wired/wireless computer or telecommunications networks, combinations of the same or the like. The communication pathways **114** in some cases may also include application programming interfaces (APIs) including, e.g., cloud service provider APIs, virtual machine management APIs, and hosted service provider APIs.

Primary Data and Exemplary Primary Storage Devices

Primary data **112** according to some embodiments is production data or other “live” data generated by the operating system and other applications **110** residing on a client computing device **102**. The primary data **112** is generally stored on the primary storage device(s) **104** and is organized via a file system supported by the client computing device **102**. For instance, the client computing device(s) **102** and corresponding applications **110** may create, access, modify, write, delete, and otherwise use primary data **112**. In some cases, some or all of the primary data **112** can be stored in cloud storage resources.

Primary data **112** is generally in the native format of the source application **110**. According to certain aspects, primary data **112** is an initial or first (e.g., created before any other copies or before at least one other copy) stored copy of data generated by the source application **110**. Primary data **112** in some cases is created substantially directly from data generated by the corresponding source applications **110**.

The primary data **112** may sometimes be referred to as a “primary copy” in the sense that it is a discrete set of data. However, the use of this term does not necessarily imply that the “primary copy” is a copy in the sense that it was copied or otherwise derived from another stored version.

The primary storage devices **104** storing the primary data **112** may be relatively fast and/or expensive (e.g., a disk drive, a hard-disk array, solid state memory, etc.). In addition, primary data **112** may be intended for relatively short term retention (e.g., several hours, days, or weeks).

According to some embodiments, the client computing device **102** can access primary data **112** from the primary storage device **104** by making conventional file system calls

via the operating system. Primary data **112** representing files may include structured data (e.g., database files), unstructured data (e.g., documents), and/or semi-structured data. Some specific examples are described below with respect to FIG. **1B**.

It can be useful in performing certain tasks to organize the primary data **112** into units of different granularities. In general, primary data **112** can include files, directories, file system volumes, data blocks, extents, or any other hierarchies or organizations of data objects. As used herein, a “data object” can refer to both (1) any file that is currently addressable by a file system or that was previously addressable by the file system (e.g., an archive file) and (2) a subset of such a file (e.g., a data block).

As will be described in further detail, it can also be useful in performing certain functions of the information management system **100** to access and modify metadata within the primary data **112**. Metadata generally includes information about data objects or characteristics associated with the data objects.

Metadata can include, without limitation, one or more of the following: the data owner (e.g., the client or user that generates the data), the last modified time (e.g., the time of the most recent modification of the data object), a data object name (e.g., a file name), a data object size (e.g., a number of bytes of data), information about the content (e.g., an indication as to the existence of a particular search term), user-supplied tags, to/from information for email (e.g., an email sender, recipient, etc.), creation date, file type (e.g., format or application type), last accessed time, application type (e.g., type of application that generated the data object), location/network (e.g., a current, past or future location of the data object and network pathways to/from the data object), geographic location (e.g., GPS coordinates), frequency of change (e.g., a period in which the data object is modified), business unit (e.g., a group or department that generates, manages or is otherwise associated with the data object), aging information (e.g., a schedule, such as a time period, in which the data object is migrated to secondary or long term storage), boot sectors, partition layouts, file location within a file folder directory structure, user permissions, owners, groups, access control lists [ACLs], system metadata (e.g., registry information), combinations of the same or the other similar information related to the data object.

In addition to metadata generated by or related to file systems and operating systems, some of the applications **110** and/or other components of the information management system **100** maintain indices of metadata for data objects, e.g., metadata associated with individual email messages. Thus, each data object may be associated with corresponding metadata. The use of metadata to perform classification and other functions is described in greater detail below.

Each of the client computing devices **102** are generally associated with and/or in communication with one or more of the primary storage devices **104** storing corresponding primary data **112**. A client computing device **102** may be considered to be “associated with” or “in communication with” a primary storage device **104** if it is capable of one or more of: routing and/or storing data to the particular primary storage device **104**, coordinating the routing and/or storing of data to the particular primary storage device **104**, retrieving data from the particular primary storage device **104**, coordinating the retrieval of data from the particular primary storage device **104**, and modifying and/or deleting data retrieved from the particular primary storage device **104**.

The primary storage devices **104** can include any of the different types of storage devices described above, or some

other kind of suitable storage device. The primary storage devices **104** may have relatively fast I/O times and/or are relatively expensive in comparison to the secondary storage devices **108**. For example, the information management system **100** may generally regularly access data and meta-

data stored on primary storage devices **104**, whereas data and metadata stored on the secondary storage devices **108** is accessed relatively less frequently. In some cases, each primary storage device **104** is dedicated to an associated client computing device **102**. For instance, a primary storage device **104** in one embodiment is a local disk drive of a corresponding client computing device **102**. In other cases, one or more primary storage devices **104** can be shared by multiple client computing devices **102**, e.g., via a network such as in a cloud storage implementation. As one example, a primary storage device **104** can be a disk array shared by a group of client computing devices **102**, such as one of the following types of disk arrays: EMC Clariion, EMC Symmetrix, EMC Celerra, Dell EqualLogic, IBM XIV, NetApp FAS, HP EVA, and HP 3PAR.

The information management system **100** may also include hosted services (not shown), which may be hosted in some cases by an entity other than the organization that employs the other components of the information management system **100**. For instance, the hosted services may be provided by various online service providers to the organization. Such service providers can provide services including social networking services, hosted email services, or hosted productivity applications or other hosted applications).

Hosted services may include software-as-a-service (SaaS), platform-as-a-service (PaaS), application service providers (ASPs), cloud services, or other mechanisms for delivering functionality via a network. As it provides services to users, each hosted service may generate additional data and metadata under management of the information management system **100**, e.g., as primary data **112**. In some cases, the hosted services may be accessed using one of the applications **110**. As an example, a hosted mail service may be accessed via browser running on a client computing device **102**. The hosted services may be implemented in a variety of computing environments. In some cases, they are implemented in an environment having a similar arrangement to the information management system **100**, where various physical and logical components are distributed over a network.

Secondary Copies and Exemplary Secondary Storage Devices

The primary data **112** stored on the primary storage devices **104** may be compromised in some cases, such as when an employee deliberately or accidentally deletes or overwrites primary data **112** during their normal course of work. Or the primary storage devices **104** can be damaged or otherwise corrupted.

For recovery and/or regulatory compliance purposes, it is therefore useful to generate copies of the primary data **112**. Accordingly, the information management system **100** includes one or more secondary storage computing devices **106** and one or more secondary storage devices **108** configured to create and store one or more secondary copies **116** of the primary data **112** and associated metadata. The secondary storage computing devices **106** and the secondary storage devices **108** may sometimes be referred to as a secondary storage subsystem **118**.

Creation of secondary copies **116** can help in search and analysis efforts and meet other information management

goals, such as: restoring data and/or metadata if an original version (e.g., of primary data **112**) is lost (e.g., by deletion, corruption, or disaster); allowing point-in-time recovery; complying with regulatory data retention and electronic discovery (e-discovery) requirements; reducing utilized storage capacity; facilitating organization and search of data; improving user access to data files across multiple computing devices and/or hosted services; and implementing data retention policies.

The client computing devices **102** access or receive primary data **112** and communicate the data, e.g., over the communication pathways **114**, for storage in the secondary storage device(s) **108**.

A secondary copy **116** can comprise a separate stored copy of application data that is derived from one or more earlier-created, stored copies (e.g., derived from primary data **112** or another secondary copy **116**). Secondary copies **116** can include point-in-time data, and may be intended for relatively long-term retention (e.g., weeks, months or years), before some or all of the data is moved to other storage or is discarded.

In some cases, a secondary copy **116** is a copy of application data created and stored subsequent to at least one other stored instance (e.g., subsequent to corresponding primary data **112** or to another secondary copy **116**), in a different storage device than at least one previous stored copy, and/or remotely from at least one previous stored copy. In some other cases, secondary copies can be stored in the same storage device as primary data **112** and/or other previously stored copies. For example, in one embodiment a disk array capable of performing hardware snapshots stores primary data **112** and creates and stores hardware snapshots of the primary data **112** as secondary copies **116**. Secondary copies **116** may be stored in relatively slow and/or low cost storage (e.g., magnetic tape). A secondary copy **116** may be stored in a backup or archive format, or in some other format different than the native source application format or other primary data format.

In some cases, secondary copies **116** are indexed so users can browse and restore at another point in time. After creation of a secondary copy **116** representative of certain primary data **112**, a pointer or other location indicia (e.g., a stub) may be placed in primary data **112**, or be otherwise associated with primary data **112** to indicate the current location on the secondary storage device(s) **108**.

Since an instance of a data object or metadata in primary data **112** may change over time as it is modified by an application **110** (or hosted service or the operating system), the information management system **100** may create and manage multiple secondary copies **116** of a particular data object or metadata, each representing the state of the data object in primary data **112** at a particular point in time. Moreover, since an instance of a data object in primary data **112** may eventually be deleted from the primary storage device **104** and the file system, the information management system **100** may continue to manage point-in-time representations of that data object, even though the instance in primary data **112** no longer exists.

For virtualized computing devices the operating system and other applications **110** of the client computing device(s) **102** may execute within or under the management of virtualization software (e.g., a VMM), and the primary storage device(s) **104** may comprise a virtual disk created on a physical storage device. The information management system **100** may create secondary copies **116** of the files or other

data objects in a virtual disk file and/or secondary copies **116** of the entire virtual disk file itself (e.g., of an entire .vmdk file).

Secondary copies **116** may be distinguished from corresponding primary data **112** in a variety of ways, some of which will now be described. First, as discussed, secondary copies **116** can be stored in a different format (e.g., backup, archive, or other non-native format) than primary data **112**. For this or other reasons, secondary copies **116** may not be directly useable by the applications **110** of the client computing device **102**, e.g., via standard system calls or otherwise without modification, processing, or other intervention by the information management system **100**.

Secondary copies **116** are also in some embodiments stored on a secondary storage device **108** that is inaccessible to the applications **110** running on the client computing devices **102** (and/or hosted services). Some secondary copies **116** may be “offline copies,” in that they are not readily available (e.g., not mounted to tape or disk). Offline copies can include copies of data that the information management system **100** can access without human intervention (e.g., tapes within an automated tape library, but not yet mounted in a drive), and copies that the information management system **100** can access only with at least some human intervention (e.g., tapes located at an offsite storage site).

The Use of Intermediate Devices for Creating Secondary Copies

Creating secondary copies can be a challenging task. For instance, there can be hundreds or thousands of client computing devices **102** continually generating large volumes of primary data **112** to be protected. Also, there can be significant overhead involved in the creation of secondary copies **116**. Moreover, secondary storage devices **108** may be special purpose components, and interacting with them can require specialized intelligence.

In some cases, the client computing devices **102** interact directly with the secondary storage device **108** to create the secondary copies **116**. However, in view of the factors described above, this approach can negatively impact the ability of the client computing devices **102** to serve the applications **110** and produce primary data **112**. Further, the client computing devices **102** may not be optimized for interaction with the secondary storage devices **108**.

Thus, in some embodiments, the information management system **100** includes one or more software and/or hardware components which generally act as intermediaries between the client computing devices **102** and the secondary storage devices **108**. In addition to off-loading certain responsibilities from the client computing devices **102**, these intermediate components can provide other benefits. For instance, as discussed further below with respect to FIG. 1D, distributing some of the work involved in creating secondary copies **116** can enhance scalability.

The intermediate components can include one or more secondary storage computing devices **106** as shown in FIG. 1A and/or one or more media agents, which can be software modules residing on corresponding secondary storage computing devices **106** (or other appropriate devices). Media agents are discussed below (e.g., with respect to FIGS. 1C-1E).

The secondary storage computing device(s) **106** can comprise any of the computing devices described above, without limitation. In some cases, the secondary storage computing device(s) **106** include specialized hardware and/or software componentry for interacting with the secondary storage devices **108**.

To create a secondary copy **116** involving the copying of data from the primary storage subsystem **117** to the secondary storage subsystem **118**, the client computing device **102** in some embodiments communicates the primary data **112** to be copied (or a processed version thereof) to the designated secondary storage computing device **106**, via the communication pathway **114**. The secondary storage computing device **106** in turn conveys the received data (or a processed version thereof) to the secondary storage device **108**. In some such configurations, the communication pathway **114** between the client computing device **102** and the secondary storage computing device **106** comprises a portion of a LAN, WAN or SAN. In other cases, at least some client computing devices **102** communicate directly with the secondary storage devices **108** (e.g., via Fibre Channel or SCSI connections). In some other cases, one or more secondary copies **116** are created from existing secondary copies, such as in the case of an auxiliary copy operation, described in greater detail below.

Exemplary Primary Data and an Exemplary Secondary Copy

FIG. 1B is a detailed view showing some specific examples of primary data stored on the primary storage device(s) **104** and secondary copy data stored on the secondary storage device(s) **108**, with other components in the system removed for the purposes of illustration. Stored on the primary storage device(s) **104** are primary data objects including word processing documents **119A-B**, spreadsheets **120**, presentation documents **122**, video files **124**, image files **126**, email mailboxes **128** (and corresponding email messages **129A-C**), html/xml or other types of markup language files **130**, databases **132** and corresponding tables or other data structures **133A-133C**).

Some or all primary data objects are associated with corresponding metadata (e.g., “Meta1-11”) which may include file system metadata and/or application specific metadata. Stored on the secondary storage device(s) **108** are secondary copy data objects **134A-C** which may include copies of or otherwise represent corresponding primary data objects and metadata.

As shown, the secondary copy data objects **134A-C** can individually represent more than one primary data object. For example, secondary copy data object **134A** represents three separate primary data objects **133C**, **122** and **129C** (represented as **133C'**, **122'** and **129C'**, respectively, and accompanied by the corresponding metadata Meta11, Meta3, and Meta8, respectively). Moreover, as indicated by the prime mark ('), a secondary copy object may store a representation of a primary data object or metadata differently than the original format, e.g., in a compressed, encrypted, deduplicated, or other modified format. Likewise, secondary data object **1346** represents primary data objects **120**, **1336**, and **119A** as **120'**, **1336'**, and **119A'**, respectively and accompanied by corresponding metadata Meta2, Meta10, and Meta1, respectively. Also, secondary data object **134C** represents primary data objects **133A**, **1196**, and **129A** as **133A'**, **1196'**, and **129A'**, respectively, accompanied by corresponding metadata Meta9, Meta5, and Meta6, respectively.

Exemplary Information Management System Architecture

The information management system **100** can incorporate a variety of different hardware and software components, which can in turn be organized with respect to one another in many different configurations, depending on the embodiment. There are critical design choices involved in specifying the functional responsibilities of the components and the

role of each component in the information management system **100**. For instance, as will be discussed, such design choices can impact performance as well as the adaptability of the information management system **100** to data growth or other changing circumstances.

FIG. 1C shows an information management system **100** designed according to these considerations and which includes: storage manager **140**, a centralized storage and/or information manager that is configured to perform certain control functions, one or more data agents **142** executing on the client computing device(s) **102** configured to process primary data **112**, and one or more media agents **144** executing on the one or more secondary storage computing devices **106** for performing tasks involving the secondary storage devices **108**. While distributing functionality amongst multiple computing devices can have certain advantages, in other contexts it can be beneficial to consolidate functionality on the same computing device. As such, in various other embodiments, one or more of the components shown in FIG. 1C as being implemented on separate computing devices are implemented on the same computing device. In one configuration, a storage manager **140**, one or more data agents **142**, and one or more media agents **144** are all implemented on the same computing device. In another embodiment, one or more data agents **142** and one or more media agents **144** are implemented on the same computing device, while the storage manager **140** is implemented on a separate computing device.

Storage Manager

As noted, the number of components in the information management system **100** and the amount of data under management can be quite large. Managing the components and data is therefore a significant task, and a task that can grow in an often unpredictable fashion as the quantity of components and data scale to meet the needs of the organization.

For these and other reasons, according to certain embodiments, responsibility for controlling the information management system **100**, or at least a significant portion of that responsibility, is allocated to the storage manager **140**.

By distributing control functionality in this manner, the storage manager **140** can be adapted independently according to changing circumstances. Moreover, a computing device for hosting the storage manager **140** can be selected to best suit the functions of the storage manager **140**. These and other advantages are described in further detail below with respect to FIG. 1D.

The storage manager **140** may be a software module or other application. In some embodiments, storage manager **140** is a computing device comprising circuitry for executing computer instructions and performs the functions described herein. The storage manager generally initiates, performs, coordinates and/or controls storage and other information management operations performed by the information management system **100**, e.g., to protect and control the primary data **112** and secondary copies **116** of data and metadata.

As shown by the dashed arrowed lines **114**, the storage manager **140** may communicate with and/or control some or all elements of the information management system **100**, such as the data agents **142** and media agents **144**. Thus, in certain embodiments, control information originates from the storage manager **140**, whereas payload data and payload metadata is generally communicated between the data agents **142** and the media agents **144** (or otherwise between the client computing device(s) **102** and the secondary storage computing device(s) **106**), e.g., at the direction of the

storage manager **140**. Control information can generally include parameters and instructions for carrying out information management operations, such as, without limitation, instructions to perform a task associated with an operation, timing information specifying when to initiate a task associated with an operation, data path information specifying what components to communicate with or access in carrying out an operation, and the like. Payload data, on the other hand, can include the actual data involved in the storage operation, such as content data written to a secondary storage device **108** in a secondary copy operation. Payload metadata can include any of the types of metadata described herein, and may be written to a storage device along with the payload content data (e.g., in the form of a header).

In other embodiments, some information management operations are controlled by other components in the information management system **100** (e.g., the media agent(s) **144** or data agent(s) **142**), instead of or in combination with the storage manager **140**.

According to certain embodiments, the storage manager **140** provides one or more of the following functions:

- initiating execution of secondary copy operations;
- managing secondary storage devices **108** and inventory/capacity of the same;
- reporting, searching, and/or classification of data in the information management system **100**;
- allocating secondary storage devices **108** for secondary storage operations;
- monitoring completion of and providing status reporting related to secondary storage operations;
- tracking age information relating to secondary copies **116**, secondary storage devices **108**, and comparing the age information against retention guidelines;
- tracking movement of data within the information management system **100**;
- tracking logical associations between components in the information management system **100**;
- protecting metadata associated with the information management system **100**; and
- implementing operations management functionality.

The storage manager **140** may maintain a database **146** (or “storage manager database **146**” or “management database **146**”) of management-related data and information management policies **148**. The database **146** may include a management index **150** (or “index **150**”) or other data structure that stores logical associations between components of the system, user preferences and/or profiles (e.g., preferences regarding encryption, compression, or deduplication of primary or secondary copy data, preferences regarding the scheduling, type, or other aspects of primary or secondary copy or other operations, mappings of particular information management users or user accounts to certain computing devices or other components, etc.), management tasks, media containerization, or other useful data. For example, the storage manager **140** may use the index **150** to track logical associations between media agents **144** and secondary storage devices **108** and/or movement of data from primary storage devices **104** to secondary storage devices **108**. For instance, the index **150** may store data associating a client computing device **102** with a particular media agent **144** and/or secondary storage device **108**, as specified in an information management policy **148** (e.g., a storage policy, which is defined in more detail below).

Administrators and other employees may be able to manually configure and initiate certain information management operations on an individual basis. But while this may be acceptable for some recovery operations or other rela-

tively less frequent tasks, it is often not workable for implementing on-going organization-wide data protection and management.

Thus, the information management system **100** may utilize information management policies **148** for specifying and executing information management operations (e.g., on an automated basis). Generally, an information management policy **148** can include a data structure or other information source that specifies a set of parameters (e.g., criteria and rules) associated with storage or other information management operations.

The storage manager database **146** may maintain the information management policies **148** and associated data, although the information management policies **148** can be stored in any appropriate location. For instance, an information management policy **148** such as a storage policy may be stored as metadata in a media agent database **152** or in a secondary storage device **108** (e.g., as an archive copy) for use in restore operations or other information management operations, depending on the embodiment. Information management policies **148** are described further below.

According to certain embodiments, the storage manager database **146** comprises a relational database (e.g., an SQL database) for tracking metadata, such as metadata associated with secondary copy operations (e.g., what client computing devices **102** and corresponding data were protected). This and other metadata may additionally be stored in other locations, such as at the secondary storage computing devices **106** or on the secondary storage devices **108**, allowing data recovery without the use of the storage manager **140**.

As shown, the storage manager **140** may include a jobs agent **156**, a user interface **158**, and a management agent **154**, all of which may be implemented as interconnected software modules or application programs.

The jobs agent **156** in some embodiments initiates, controls, and/or monitors the status of some or all storage or other information management operations previously performed, currently being performed, or scheduled to be performed by the information management system **100**. For instance, the jobs agent **156** may access information management policies **148** to determine when and how to initiate and control secondary copy and other information management operations, as will be discussed further.

The user interface **158** may include information processing and display software, such as a graphical user interface (“GUI”), an application program interface (“API”), or other interactive interface(s) through which users and system processes can retrieve information about the status of information management operations (e.g., storage operations) or issue instructions to the information management system **100** and its constituent components.

Via the user interface **158**, users may optionally issue instructions to the components in the information management system **100** regarding performance of storage and recovery operations. For example, a user may modify a schedule concerning the number of pending secondary copy operations. As another example, a user may employ the GUI to view the status of pending storage operations or to monitor the status of certain components in the information management system **100** (e.g., the amount of capacity left in a storage device).

An information management “cell” may generally include a logical and/or physical grouping of a combination of hardware and software components associated with performing information management operations on electronic data, typically one storage manager **140** and at least one

client computing device **102** (comprising data agent(s) **142**) and at least one media agent **144**. For instance, the components shown in FIG. **1C** may together form an information management cell. Multiple cells may be organized hierarchically. With this configuration, cells may inherit properties from hierarchically superior cells or be controlled by other cells in the hierarchy (automatically or otherwise). Alternatively, in some embodiments, cells may inherit or otherwise be associated with information management policies, preferences, information management metrics, or other properties or characteristics according to their relative position in a hierarchy of cells. Cells may also be delineated and/or organized hierarchically according to function, geography, architectural considerations, or other factors useful or desirable in performing information management operations. A first cell may represent a geographic segment of an enterprise, such as a Chicago office, and a second cell may represent a different geographic segment, such as a New York office. Other cells may represent departments within a particular office. Where delineated by function, a first cell may perform one or more first types of information management operations (e.g., one or more first types of secondary or other copies), and a second cell may perform one or more second types of information management operations (e.g., one or more second types of secondary or other copies).

The storage manager **140** may also track information that permits it to select, designate, or otherwise identify content indices, deduplication databases, or similar databases or resources or data sets within its information management cell (or another cell) to be searched in response to certain queries. Such queries may be entered by the user via interaction with the user interface **158**. In general, the management agent **154** allows multiple information management cells to communicate with one another. For example, the information management system **100** in some cases may be one information management cell of a network of multiple cells adjacent to one another or otherwise logically related in a WAN or LAN. With this arrangement, the cells may be connected to one another through respective management agents **154**.

For instance, the management agent **154** can provide the storage manager **140** with the ability to communicate with other components within the information management system **100** (and/or other cells within a larger information management system) via network protocols and application programming interfaces (“APIs”) including, e.g., HTTP, HTTPS, FTP, REST, virtualization software APIs, cloud service provider APIs, and hosted service provider APIs. Inter-cell communication and hierarchy is described in greater detail in U.S. Pat. Nos. 7,747,579 and 7,343,453, which are incorporated by reference herein.

Data Agents

As discussed, a variety of different types of applications **110** can reside on a given client computing device **102**, including operating systems, database applications, e-mail applications, and virtual machines, just to name a few. And, as part of the process of creating and restoring secondary copies **116**, the client computing devices **102** may be tasked with processing and preparing the primary data **112** from these various different applications **110**. Moreover, the nature of the processing/preparation can differ across clients and application types, e.g., due to inherent structural and formatting differences between applications **110**.

The one or more data agent(s) **142** are therefore advantageously configured in some embodiments to assist in the performance of information management operations based

on the type of data that is being protected, at a client-specific and/or application-specific level.

The data agent **142** may be a software module or component that is generally responsible for managing, initiating, or otherwise assisting in the performance of information management operations. For instance, the data agent **142** may take part in performing data storage operations such as the copying, archiving, migrating, replicating of primary data **112** stored in the primary storage device(s) **104**. The data agent **142** may receive control information from the storage manager **140**, such as commands to transfer copies of data objects, metadata, and other payload data to the media agents **144**.

In some embodiments, a data agent **142** may be distributed between the client computing device **102** and storage manager **140** (and any other intermediate components) or may be deployed from a remote location or its functions approximated by a remote process that performs some or all of the functions of data agent **142**. In addition, a data agent **142** may perform some functions provided by a media agent **144**, or may perform other functions such as encryption and deduplication.

As indicated, each data agent **142** may be specialized for a particular application **110**, and the system can employ multiple application-specific data agents **142**, each of which may perform information management operations (e.g., perform backup, migration, and data recovery) associated with a different application **110**. For instance, different individual data agents **142** may be designed to handle Microsoft Exchange data, Lotus Notes data, Microsoft Windows file system data, Microsoft Active Directory Objects data, SQL Server data, SharePoint data, Oracle database data, SAP database data, virtual machines and/or associated data, and other types of data.

A file system data agent, for example, may handle data files and/or other file system information. If a client computing device **102** has two or more types of data, one data agent **142** may be used for each data type to copy, archive, migrate, and restore the client computing device **102** data. For example, to backup, migrate, and restore all of the data on a Microsoft Exchange server, the client computing device **102** may use one Microsoft Exchange Mailbox data agent **142** to backup the Exchange mailboxes, one Microsoft Exchange Database data agent **142** to backup the Exchange databases, one Microsoft Exchange Public Folder data agent **142** to backup the Exchange Public Folders, and one Microsoft Windows File System data agent **142** to backup the file system of the client computing device **102**. In such embodiments, these data agents **142** may be treated as four separate data agents **142** even though they reside on the same client computing device **102**.

Other embodiments may employ one or more generic data agents **142** that can handle and process data from two or more different applications **110**, or that can handle and process multiple data types, instead of or in addition to using specialized data agents **142**. For example, one generic data agent **142** may be used to back up, migrate and restore Microsoft Exchange Mailbox data and Microsoft Exchange Database data while another generic data agent may handle Microsoft Exchange Public Folder data and Microsoft Windows File System data.

Each data agent **142** may be configured to access data and/or metadata stored in the primary storage device(s) **104** associated with the data agent **142** and process the data as appropriate. For example, during a secondary copy operation, the data agent **142** may arrange or assemble the data and metadata into one or more files having a certain format

(e.g., a particular backup or archive format) before transferring the file(s) to a media agent **144** or other component. The file(s) may include a list of files or other metadata. Each data agent **142** can also assist in restoring data or metadata to primary storage devices **104** from a secondary copy **116**. For instance, the data agent **142** may operate in conjunction with the storage manager **140** and one or more of the media agents **144** to restore data from secondary storage device(s) **108**.

Media Agents

As indicated above with respect to FIG. **1A**, off-loading certain responsibilities from the client computing devices **102** to intermediate components such as the media agent(s) **144** can provide a number of benefits including improved client computing device **102** operation, faster secondary copy operation performance, and enhanced scalability. As one specific example which will be discussed below in further detail, the media agent **144** can act as a local cache of copied data and/or metadata that it has stored to the secondary storage device(s) **108**, providing improved restore capabilities.

Generally speaking, a media agent **144** may be implemented as a software module that manages, coordinates, and facilitates the transmission of data, as directed by the storage manager **140**, between a client computing device **102** and one or more secondary storage devices **108**. Whereas the storage manager **140** controls the operation of the information management system **100**, the media agent **144** generally provides a portal to secondary storage devices **108**. For instance, other components in the system interact with the media agents **144** to gain access to data stored on the secondary storage devices **108**, whether it be for the purposes of reading, writing, modifying, or deleting data. Moreover, as will be described further, media agents **144** can generate and store information relating to characteristics of the stored data and/or metadata, or can generate and store other types of information that generally provides insight into the contents of the secondary storage devices **108**.

Media agents **144** can comprise separate nodes in the information management system **100** (e.g., nodes that are separate from the client computing devices **102**, storage manager **140**, and/or secondary storage devices **108**). In general, a node within the information management system **100** can be a logically and/or physically separate component, and in some cases is a component that is individually addressable or otherwise identifiable. In addition, each media agent **144** may reside on a dedicated secondary storage computing device **106** in some cases, while in other embodiments a plurality of media agents **144** reside on the same secondary storage computing device **106**.

A media agent **144** (and corresponding media agent database **152**) may be considered to be "associated with" a particular secondary storage device **108** if that media agent **144** is capable of one or more of: routing and/or storing data to the particular secondary storage device **108**, coordinating the routing and/or storing of data to the particular secondary storage device **108**, retrieving data from the particular secondary storage device **108**, coordinating the retrieval of data from a particular secondary storage device **108**, and modifying and/or deleting data retrieved from the particular secondary storage device **108**.

While media agent(s) **144** are generally associated with one or more secondary storage devices **108**, one or more media agents **144** in certain embodiments are physically separate from the secondary storage devices **108**. For instance, the media agents **144** may reside on secondary storage computing devices **106** having different housings or

packages than the secondary storage devices **108**. In one example, a media agent **144** resides on a first server computer and is in communication with a secondary storage device(s) **108** residing in a separate, rack-mounted RAID-based system.

Where the information management system **100** includes multiple media agents **144** (FIG. 1D), a first media agent **144** may provide failover functionality for a second, failed media agent **144**. In addition, media agents **144** can be dynamically selected for storage operations to provide load balancing. Failover and load balancing are described in greater detail below.

In operation, a media agent **144** associated with a particular secondary storage device **108** may instruct the secondary storage device **108** to perform an information management operation. For instance, a media agent **144** may instruct a tape library to use a robotic arm or other retrieval means to load or eject a certain storage media, and to subsequently archive, migrate, or retrieve data to or from that media, e.g., for the purpose of restoring the data to a client computing device **102**. As another example, a secondary storage device **108** may include an array of hard disk drives or solid state drives organized in a RAID configuration, and the media agent **144** may forward a logical unit number (LUN) and other appropriate information to the array, which uses the received information to execute the desired storage operation. The media agent **144** may communicate with a secondary storage device **108** via a suitable communications link, such as a SCSI or Fiber Channel link.

As shown, each media agent **144** may maintain an associated media agent database **152**. The media agent database **152** may be stored in a disk or other storage device (not shown) that is local to the secondary storage computing device **106** on which the media agent **144** resides. In other cases, the media agent database **152** is stored remotely from the secondary storage computing device **106**.

The media agent database **152** can include, among other things, an index **153** including data generated during secondary copy operations and other storage or information management operations. The index **153** provides a media agent **144** or other component with a fast and efficient mechanism for locating secondary copies **116** or other data stored in the secondary storage devices **108**. In some cases, the index **153** does not form a part of and is instead separate from the media agent database **152**.

A media agent index **153** or other data structure associated with the particular media agent **144** may include information about the stored data. For instance, for each secondary copy **116**, the index **153** may include metadata such as a list of the data objects (e.g., files/subdirectories, database objects, mailbox objects, etc.), a path to the secondary copy **116** on the corresponding secondary storage device **108**, location information indicating where the data objects are stored in the secondary storage device **108**, when the data objects were created or modified, etc. Thus, the index **153** includes metadata associated with the secondary copies **116** that is readily available for use in storage operations and other activities without having to be first retrieved from the secondary storage device **108**. In yet further embodiments, some or all of the data in the index **153** may instead or additionally be stored along with the data in a secondary storage device **108**, e.g., with a copy of the index **153**. In some embodiments, the secondary storage devices **108** can include sufficient information to perform a “bare metal restore”, where the operating system of a failed client computing device **102** or other restore target is automatically rebuilt as part of a restore operation.

Because the index **153** maintained in the media agent database **152** may operate as a cache, it can also be referred to as “an index cache.” In such cases, information stored in the index cache **153** typically comprises data that reflects certain particulars about storage operations that have occurred relatively recently. After some triggering event, such as after a certain period of time elapses, or the index cache **153** reaches a particular size, the index cache **153** may be copied or migrated to a secondary storage device(s) **108**. This information may need to be retrieved and uploaded back into the index cache **153** or otherwise restored to a media agent **144** to facilitate retrieval of data from the secondary storage device(s) **108**. In some embodiments, the cached information may include format or containerization information related to archives or other files stored on the storage device(s) **108**. In this manner, the index cache **153** allows for accelerated restores.

In some alternative embodiments the media agent **144** generally acts as a coordinator or facilitator of storage operations between client computing devices **102** and corresponding secondary storage devices **108**, but does not actually write the data to the secondary storage device **108**. For instance, the storage manager **140** (or the media agent **144**) may instruct a client computing device **102** and secondary storage device **108** to communicate with one another directly. In such a case the client computing device **102** transmits the data directly or via one or more intermediary components to the secondary storage device **108** according to the received instructions, and vice versa. In some such cases, the media agent **144** may still receive, process, and/or maintain metadata related to the storage operations. Moreover, in these embodiments, the payload data can flow through the media agent **144** for the purposes of populating the index cache **153** maintained in the media agent database **152**, but not for writing to the secondary storage device **108**.

The media agent **144** and/or other components such as the storage manager **140** may in some cases incorporate additional functionality, such as data classification, content indexing, deduplication, encryption, compression, and the like. Further details regarding these and other functions are described below.

Distributed, Scalable Architecture

As described, certain functions of the information management system **100** can be distributed amongst various physical and/or logical components in the system. For instance, one or more of the storage manager **140**, data agents **142**, and media agents **144** may reside on computing devices that are physically separate from one another. This architecture can provide a number of benefits.

For instance, hardware and software design choices for each distributed component can be targeted to suit its particular function. The secondary computing devices **106** on which the media agents **144** reside can be tailored for interaction with associated secondary storage devices **108** and provide fast index cache operation, among other specific tasks. Similarly, the client computing device(s) **102** can be selected to effectively service the applications **110** residing thereon, in order to efficiently produce and store primary data **112**.

Moreover, in some cases, one or more of the individual components in the information management system **100** can be distributed to multiple, separate computing devices. As one example, for large file systems where the amount of data stored in the database **146** is relatively large, the database **146** may be migrated to or otherwise reside on a specialized database server (e.g., an SQL server) separate from a server that implements the other functions of the storage manager

140. This configuration can provide added protection because the database 146 can be protected with standard database utilities (e.g., SQL log shipping or database replication) independent from other functions of the storage manager 140. The database 146 can be efficiently replicated to a remote site for use in the event of a disaster or other data loss incident at the primary site. Or the database 146 can be replicated to another computing device within the same site, such as to a higher performance machine in the event that a storage manager host device can no longer service the needs of a growing information management system 100.

The distributed architecture also provides both scalability and efficient component utilization. FIG. 1D shows an embodiment of the information management system 100 including a plurality of client computing devices 102 and associated data agents 142 as well as a plurality of secondary storage computing devices 106 and associated media agents 144.

Additional components can be added or subtracted based on the evolving needs of the information management system 100. For instance, depending on where bottlenecks are identified, administrators can add additional client computing devices 102, secondary storage computing devices 106 (and corresponding media agents 144), and/or secondary storage devices 108. Moreover, where multiple fungible components are available, load balancing can be implemented to dynamically address identified bottlenecks. As an example, the storage manager 140 may dynamically select which media agents 144 and/or secondary storage devices 108 to use for storage operations based on a processing load analysis of the media agents 144 and/or secondary storage devices 108, respectively.

Moreover, each client computing device 102 in some embodiments can communicate with, among other components, any of the media agents 144, e.g., as directed by the storage manager 140. And each media agent 144 may be able to communicate with, among other components, any of the secondary storage devices 108, e.g., as directed by the storage manager 140. Thus, operations can be routed to the secondary storage devices 108 in a dynamic and highly flexible manner, to provide load balancing, failover, and the like. Further examples of scalable systems capable of dynamic storage operations, and of systems capable of performing load balancing and fail over are provided in U.S. Pat. No. 7,246,207, which is incorporated by reference herein.

In alternative configurations, certain components are not distributed and may instead reside and execute on the same computing device. For example, in some embodiments one or more data agents 142 and the storage manager 140 reside on the same client computing device 102. In another embodiment, one or more data agents 142 and one or more media agents 144 reside on a single computing device.

Exemplary Types of Information Management Operations

In order to protect and leverage stored data, the information management system 100 can be configured to perform a variety of information management operations. As will be described, these operations can generally include secondary copy and other data movement operations, processing and data manipulation operations, analysis, reporting, and management operations. The operations described herein may be performed on any type of computing platform, e.g., between two computers connected via a LAN, to a mobile client telecommunications device connected to a server via a WLAN, to any manner of client device coupled to a cloud storage target.

Data Movement Operations

Data movement operations according to certain embodiments are generally operations that involve the copying or migration of data (e.g., payload data) between different locations in the information management system 100 in an original/native and/or one or more different formats. For example, data movement operations can include operations in which stored data is copied, migrated, or otherwise transferred from one or more first storage devices to one or more second storage devices, such as from primary storage device(s) 104 to secondary storage device(s) 108, from secondary storage device(s) 108 to different secondary storage device(s) 108, from secondary storage devices 108 to primary storage devices 104, or from primary storage device(s) 104 to different primary storage device(s) 104.

Data movement operations can include by way of example, backup operations, archive operations, information lifecycle management operations such as hierarchical storage management operations, replication operations (e.g., continuous data replication operations), snapshot operations, deduplication or single-instancing operations, auxiliary copy operations, and the like. As will be discussed, some of these operations involve the copying, migration or other movement of data, without actually creating multiple, distinct copies. Nonetheless, some or all of these operations are referred to as “copy” operations for simplicity.

Backup Operations

A backup operation creates a copy of a version of data (e.g., one or more files or other data units) in primary data 112 at a particular point in time. Each subsequent backup copy may be maintained independently of the first. Further, a backup copy in some embodiments is generally stored in a form that is different than the native format, e.g., a backup format. This can be in contrast to the version in primary data 112 from which the backup copy is derived, and which may instead be stored in a native format of the source application(s) 110. In various cases, backup copies can be stored in a format in which the data is compressed, encrypted, deduplicated, and/or otherwise modified from the original application format. For example, a backup copy may be stored in a backup format that facilitates compression and/or efficient long-term storage.

Backup copies can have relatively long retention periods as compared to primary data 112, and may be stored on media with slower retrieval times than primary data 112 and certain other types of secondary copies 116. On the other hand, backups may have relatively shorter retention periods than some other types of secondary copies 116, such as archive copies (described below). Backups may sometimes be stored at an offsite location.

Backup operations can include full, synthetic or incremental backups. A full backup in some embodiments is generally a complete image of the data to be protected. However, because full backup copies can consume a relatively large amount of storage, it can be useful to use a full backup copy as a baseline and only store changes relative to the full backup copy for subsequent backup copies.

For instance, a differential backup operation (or cumulative incremental backup operation) tracks and stores changes that have occurred since the last full backup. Differential backups can grow quickly in size, but can provide relatively efficient restore times because a restore can be completed in some cases using only the full backup copy and the latest differential copy.

An incremental backup operation generally tracks and stores changes since the most recent backup copy of any type, which can generally reduce storage utilization. In some

cases, however, restore times can be relatively long in comparison to full or differential backups because completing a restore operation may involve accessing a full backup in addition to multiple incremental backups.

Any of the above types of backup operations can be at the volume-level, file-level, or block-level. Volume level backup operations generally involve the copying of a data volume (e.g., a logical disk or partition) as a whole. In a file-level backup, the information management system **100** may generally track changes to individual files at the file-level, and includes copies of files in the backup copy. In the case of a block-level backup, files are broken into constituent blocks, and changes are tracked at the block-level. Upon restore, the information management system **100** reassembles the blocks into files in a transparent fashion.

Far less data may actually be transferred and copied to the secondary storage devices **108** during a file-level copy than a volume-level copy. Likewise, a block-level copy may involve the transfer of less data than a file-level copy, resulting in faster execution times. However, restoring a relatively higher-granularity copy can result in longer restore times. For instance, when restoring a block-level copy, the process of locating constituent blocks can sometimes result in longer restore times as compared to file-level backups. Similar to backup operations, the other types of secondary copy operations described herein can also be implemented at either the volume-level, file-level, or block-level.

Archive Operations

Because backup operations generally involve maintaining a version of the copied data in primary data **112** and also maintaining backup copies in secondary storage device(s) **108**, they can consume significant storage capacity. To help reduce storage consumption, an archive operation according to certain embodiments creates a secondary copy **116** by both copying and removing source data. Or, seen another way, archive operations can involve moving some or all of the source data to the archive destination. Thus, data satisfying criteria for removal (e.g., data of a threshold age or size) from the source copy may be removed from source storage. Archive copies are sometimes stored in an archive format or other non-native application format. The source data may be primary data **112** or a secondary copy **116**, depending on the situation. As with backup copies, archive copies can be stored in a format in which the data is compressed, encrypted, deduplicated, and/or otherwise modified from the original application format.

In addition, archive copies may be retained for relatively long periods of time (e.g., years) and, in some cases, are never deleted. Archive copies are generally retained for longer periods of time than backup copies, for example. In certain embodiments, archive copies may be made and kept for extended periods in order to meet compliance regulations.

Moreover, when primary data **112** is archived, in some cases the archived primary data **112** or a portion thereof is deleted when creating the archive copy. Thus, archiving can serve the purpose of freeing up space in the primary storage device(s) **104**. Similarly, when a secondary copy **116** is archived, the secondary copy **116** may be deleted, and an archive copy can therefore serve the purpose of freeing up space in secondary storage device(s) **108**. In contrast, source copies often remain intact when creating backup copies. Examples of compatible data archiving operations are provided in U.S. Pat. No. 7,107,298, which is incorporated by reference herein.

Snapshot Operations

Snapshot operations can provide a relatively lightweight, efficient mechanism for protecting data. From an end-user viewpoint, a snapshot may be thought of as an “instant” image of the primary data **112** at a given point in time, and may include state and/or status information relative to an application that creates/manages the data. In one embodiment, a snapshot may generally capture the directory structure of an object in primary data **112** such as a file or volume or other data set at a particular moment in time and may also preserve file attributes and contents. A snapshot in some cases is created relatively quickly, e.g., substantially instantly, using a minimum amount of file space, but may still function as a conventional file system backup.

A “hardware snapshot” (or “hardware-based snapshot”) operation can be a snapshot operation where a target storage device (e.g., a primary storage device **104** or a secondary storage device **108**) performs the snapshot operation in a self-contained fashion, substantially independently, using hardware, firmware and/or software residing on the storage device itself. For instance, the storage device may be capable of performing snapshot operations upon request, generally without intervention or oversight from any of the other components in the information management system **100**. In this manner, hardware snapshots can off-load other components of information management system **100** from processing involved in snapshot creation and management.

A “software snapshot” (or “software-based snapshot”) operation, on the other hand, can be a snapshot operation in which one or more other components in information management system **100** (e.g., client computing devices **102**, data agents **142**, etc.) implement a software layer that manages the snapshot operation via interaction with the target storage device. For instance, the component implementing the snapshot management software layer may derive a set of pointers and/or data that represents the snapshot. The snapshot management software layer may then transmit the same to the target storage device, along with appropriate instructions for writing the snapshot.

Some types of snapshots do not actually create another physical copy of all the data as it existed at the particular point in time, but may simply create pointers that are able to map files and directories to specific memory locations (e.g., to specific disk blocks) where the data resides, as it existed at the particular point in time. For example, a snapshot copy may include a set of pointers derived from the file system or an application. In some other cases, the snapshot may be created at the block-level, such as where creation of the snapshot occurs without awareness of the file system. Each pointer points to a respective stored data block, so that collectively, the set of pointers reflect the storage location and state of the data object (e.g., file(s) or volume(s) or data set(s)) at a particular point in time when the snapshot copy was created.

Once a snapshot has been taken, subsequent changes to the file system typically do not overwrite the blocks in use at the time of the snapshot. Therefore, the initial snapshot may use only a small amount of disk space needed to record a mapping or other data structure representing or otherwise tracking the blocks that correspond to the current state of the file system. Additional disk space is usually required only when files and directories are actually later modified. Furthermore, when files are modified, typically only the pointers which map to blocks are copied, not the blocks themselves. In some embodiments, for example in the case of “copy-on-write” snapshots, when a block changes in pri-

mary storage, the block is copied to secondary storage or cached in primary storage before the block is overwritten in primary storage, and the pointer to that block changed to reflect the new location of that block. The snapshot mapping of file system data may also be updated to reflect the changed block(s) at that particular point in time. In some other cases, a snapshot includes a full physical copy of all or substantially all of the data represented by the snapshot. Further examples of snapshot operations are provided in U.S. Pat. No. 7,529,782, which is incorporated by reference herein.

A snapshot copy in many cases can be made quickly and without significantly impacting primary computing resources because large amounts of data need not be copied or moved. In some embodiments, a snapshot may exist as a virtual file system, parallel to the actual file system. Users in some cases gain read-only access to the record of files and directories of the snapshot. By electing to restore primary data **112** from a snapshot taken at a given point in time, users may also return the current file system to the state of the file system that existed when the snapshot was taken.

Replication Operations

Another type of secondary copy operation is a replication operation. Some types of secondary copies **116** are used to periodically capture images of primary data **112** at particular points in time (e.g., backups, archives, and snapshots). However, it can also be useful for recovery purposes to protect primary data **112** in a more continuous fashion, by replicating the primary data **112** substantially as changes occur. In some cases a replication copy can be a mirror copy, for instance, where changes made to primary data **112** are mirrored or substantially immediately copied to another location (e.g., to secondary storage device(s) **108**). By copying each write operation to the replication copy, two storage systems are kept synchronized or substantially synchronized so that they are virtually identical at approximately the same time. Where entire disk volumes are mirrored, however, mirroring can require significant amount of storage space and utilizes a large amount of processing resources.

According to some embodiments storage operations are performed on replicated data that represents a recoverable state, or “known good state” of a particular application running on the source system. For instance, in certain embodiments, known good replication copies may be viewed as copies of primary data **112**. This feature allows the system to directly access, copy, restore, backup or otherwise manipulate the replication copies as if the data was the “live”, primary data **112**. This can reduce access time, storage utilization, and impact on source applications **110**, among other benefits.

Based on known good state information, the information management system **100** can replicate sections of application data that represent a recoverable state rather than rote copying of blocks of data. Examples of compatible replication operations (e.g., continuous data replication) are provided in U.S. Pat. No. 7,617,262, which is incorporated by reference herein.

Deduplication/Single-Instancing Operations

Another type of data movement operation is deduplication or single-instance storage, which is useful to reduce the amount of data within the system. For instance, some or all of the above-described secondary storage operations can involve deduplication in some fashion. New data is read, broken down into portions (e.g., sub-file level blocks, files, etc.) of a selected granularity, compared with blocks that are

already stored, and only the new blocks are stored. Blocks that already exist are represented as pointers to the already stored data.

In order to streamline the comparison process, the information management system **100** may calculate and/or store signatures (e.g., hashes or cryptographically unique IDs) corresponding to the individual data blocks in a database and compare the signatures instead of comparing entire data blocks. In some cases, only a single instance of each element is stored, and deduplication operations may therefore be referred to interchangeably as “single-instancing” operations. Depending on the implementation, however, deduplication or single-instancing operations can store more than one instance of certain data blocks, but nonetheless significantly reduce data redundancy.

Depending on the embodiment, deduplication blocks can be of fixed or variable length. Using variable length blocks can provide enhanced deduplication by responding to changes in the data stream, but can involve complex processing. In some cases, the information management system **100** utilizes a technique for dynamically aligning deduplication blocks (e.g., fixed-length blocks) based on changing content in the data stream, as described in U.S. Pat. No. 8,364,652, which is incorporated by reference herein.

The information management system **100** can perform deduplication in a variety of manners at a variety of locations in the information management system **100**. For instance, in some embodiments, the information management system **100** implements “target-side” deduplication by deduplicating data (e.g., secondary copies **116**) stored in the secondary storage devices **108**. In some such cases, the media agents **144** are generally configured to manage the deduplication process. For instance, one or more of the media agents **144** maintain a corresponding deduplication database that stores deduplication information (e.g., data-block signatures). Examples of such a configuration are provided in U.S. Pat. Pub. No. 2012/0150826, which is incorporated by reference herein. Instead of or in combination with “target-side” deduplication, deduplication can also be performed on the “source-side” (or “client-side”), e.g., to reduce the amount of traffic between the media agents **144** and the client computing device(s) **102** and/or reduce redundant data stored in the primary storage devices **104**. According to various implementations, one or more of the storage devices of the target-side, source-side, or client-side of an operation can be cloud-based storage devices. Thus, the target-side, source-side, and/or client-side deduplication can be cloud-based deduplication. In particular, as discussed previously, the storage manager **140** may communicate with other components within the information management system **100** via network protocols and cloud service provider APIs to facilitate cloud-based deduplication/single instancing. Examples of such deduplication techniques are provided in U.S. Pat. Pub. No. 2012/0150818, which is incorporated by reference herein. Some other compatible deduplication/single instancing techniques are described in U.S. Pat. Pub. Nos. 2006/0224846 and 2009/0319534, which are incorporated by reference herein.

Information Lifecycle Management and Hierarchical Storage Management Operations

In some embodiments, files and other data over their lifetime move from more expensive, quick access storage to less expensive, slower access storage. Operations associated with moving data through various tiers of storage are sometimes referred to as information lifecycle management (ILM) operations.

One type of ILM operation is a hierarchical storage management (HSM) operation. A HSM operation is generally an operation for automatically moving data between classes of storage devices, such as between high-cost and low-cost storage devices. For instance, an HSM operation may involve movement of data from primary storage devices **104** to secondary storage devices **108**, or between tiers of secondary storage devices **108**. With each tier, the storage devices may be progressively relatively cheaper, have relatively slower access/restore times, etc. For example, movement of data between tiers may occur as data becomes less important over time.

In some embodiments, an HSM operation is similar to an archive operation in that creating an HSM copy may (though not always) involve deleting some of the source data, e.g., according to one or more criteria related to the source data. For example, an HSM copy may include data from primary data **112** or a secondary copy **116** that is larger than a given size threshold or older than a given age threshold and that is stored in a backup format.

Often, and unlike some types of archive copies, HSM data that is removed or aged from the source copy is replaced by a logical reference pointer or stub. The reference pointer or stub can be stored in the primary storage device **104** (or other source storage device, such as a secondary storage device **108**) to replace the deleted data in primary data **112** (or other source copy) and to point to or otherwise indicate the new location in a secondary storage device **108**.

According to one example, files are generally moved between higher and lower cost storage depending on how often the files are accessed. When a user requests access to the HSM data that has been removed or migrated, the information management system **100** uses the stub to locate the data and often make recovery of the data appear transparent, even though the HSM data may be stored at a location different from the remaining source data. In this manner, the data appears to the user (e.g., in file system browsing windows and the like) as if it still resides in the source location (e.g., in a primary storage device **104**). The stub may also include some metadata associated with the corresponding data, so that a file system and/or application can provide some information about the data object and/or a limited-functionality version (e.g., a preview) of the data object.

An HSM copy may be stored in a format other than the native application format (e.g., where the data is compressed, encrypted, deduplicated, and/or otherwise modified from the original application format). In some cases, copies which involve the removal of data from source storage and the maintenance of stub or other logical reference information on source storage may be referred to generally as "on-line archive copies". On the other hand, copies which involve the removal of data from source storage without the maintenance of stub or other logical reference information on source storage may be referred to as "off-line archive copies". Examples of HSM and ILM techniques are provided in U.S. Pat. No. 7,343,453, which is incorporated by reference herein.

Auxiliary Copy and Disaster Recovery Operations

An auxiliary copy is generally a copy operation in which a copy is created of an existing secondary copy **116**. For instance, an initial secondary copy **116** may be generated using or otherwise be derived from primary data **112** (or other data residing in the secondary storage subsystem **118**), whereas an auxiliary copy is generated from the initial secondary copy **116**. Auxiliary copies can be used to create additional standby copies of data and may reside on different

secondary storage devices **108** than the initial secondary copies **116**. Thus, auxiliary copies can be used for recovery purposes if initial secondary copies **116** become unavailable. Exemplary compatible auxiliary copy techniques are described in further detail in U.S. Pat. No. 8,230,195, which is incorporated by reference herein.

The information management system **100** may also perform disaster recovery operations that make or retain disaster recovery copies, often as secondary, high-availability disk copies. The information management system **100** may create secondary disk copies and store the copies at disaster recovery locations using auxiliary copy or replication operations, such as continuous data replication technologies. Depending on the particular data protection goals, disaster recovery locations can be remote from the client computing devices **102** and primary storage devices **104**, remote from some or all of the secondary storage devices **108**, or both.

Data Analysis, Reporting, and Management Operations

Data analysis, reporting, and management operations can be different than data movement operations in that they do not necessarily involve the copying, migration or other transfer of data (e.g., primary data **112** or secondary copies **116**) between different locations in the system. For instance, data analysis operations may involve processing (e.g., offline processing) or modification of already stored primary data **112** and/or secondary copies **116**. However, in some embodiments data analysis operations are performed in conjunction with data movement operations. Some data analysis operations include content indexing operations and classification operations which can be useful in leveraging the data under management to provide enhanced search and other features. Other data analysis operations such as compression and encryption can provide data reduction and security benefits, respectively.

Classification Operations/Content Indexing

In some embodiments, the information management system **100** analyzes and indexes characteristics, content, and metadata associated with the data stored within the primary data **112** and/or secondary copies **116**, providing enhanced search and management capabilities for data discovery and other purposes. The content indexing can be used to identify files or other data objects having pre-defined content (e.g., user-defined keywords or phrases, other keywords/phrases that are not defined by a user, etc.), and/or metadata (e.g., email metadata such as "to", "from", "cc", "bcc", attachment name, received time, etc.).

The information management system **100** generally organizes and catalogues the results in a content index, which may be stored within the media agent database **152**, for example. The content index can also include the storage locations of (or pointer references to) the indexed data in the primary data **112** or secondary copies **116**, as appropriate. The results may also be stored, in the form of a content index database or otherwise, elsewhere in the information management system **100** (e.g., in the primary storage devices **104**, or in the secondary storage device **108**). Such index data provides the storage manager **140** or another component with an efficient mechanism for locating primary data **112** and/or secondary copies **116** of data objects that match particular criteria.

For instance, search criteria can be specified by a user through user interface **158** of the storage manager **140**. In some cases, the information management system **100** analyzes data and/or metadata in secondary copies **116** to create an "off-line" content index, without significantly impacting the performance of the client computing devices **102**. Depending on the embodiment, the system can also imple-

ment “on-line” content indexing, e.g., of primary data **112**. Examples of compatible content indexing techniques are provided in U.S. Pat. No. 8,170,995, which is incorporated by reference herein.

In order to further leverage the data stored in the information management system **100** to perform these and other tasks, one or more components can be configured to scan data and/or associated metadata for classification purposes to populate a database (or other data structure) of information (which can be referred to as a “data classification database” or a “metabase”). Depending on the embodiment, the data classification database(s) can be organized in a variety of different ways, including centralization, logical sub-divisions, and/or physical sub-divisions. For instance, one or more centralized data classification databases may be associated with different subsystems or tiers within the information management system **100**. As an example, there may be a first centralized metabase associated with the primary storage subsystem **117** and a second centralized metabase associated with the secondary storage subsystem **118**. In other cases, there may be one or more metabases associated with individual components. For instance, there may be a dedicated metabase associated with some or all of the client computing devices **102** and/or media agents **144**. In some embodiments, a data classification database may reside as one or more data structures within management database **146**, or may be otherwise associated with storage manager **140**.

In some cases, the metabase(s) may be included in separate database(s) and/or on separate storage device(s) from primary data **112** and/or secondary copies **116**, such that operations related to the metabase do not significantly impact performance on other components in the information management system **100**. In other cases, the metabase(s) may be stored along with primary data **112** and/or secondary copies **116**. Files or other data objects can be associated with identifiers (e.g., tag entries, etc.) in the media agent **144** (or other indices) to facilitate searches of stored data objects. Among a number of other benefits, the metabase can also allow efficient, automatic identification of files or other data objects to associate with secondary copy or other information management operations (e.g., in lieu of scanning an entire file system). Examples of compatible metabases and data classification operations are provided in U.S. Pat. Nos. 8,229,954 and 7,747,579, which are incorporated by reference herein.

Encryption Operations

The information management system **100** in some cases is configured to process data (e.g., files or other data objects, secondary copies **116**, etc.), according to an appropriate encryption algorithm (e.g., Blowfish, Advanced Encryption Standard [AES], Triple Data Encryption Standard [3-DES], etc.) to limit access and provide data security in the information management system **100**.

The information management system **100** in some cases encrypts the data at the client level, such that the client computing devices **102** (e.g., the data agents **142**) encrypt the data prior to forwarding the data to other components, e.g., before sending the data to media agents **144** during a secondary copy operation. In such cases, the client computing device **102** may maintain or have access to an encryption key or passphrase for decrypting the data upon restore. Encryption can also occur when creating copies of secondary copies, e.g., when creating auxiliary copies or archive copies. In yet further embodiments, the secondary storage devices **108** can implement built-in, high performance hardware encryption.

Management and Reporting Operations

Certain embodiments leverage the integrated, ubiquitous nature of the information management system **100** to provide useful system-wide management and reporting functions. Examples of some compatible management and reporting techniques are provided in U.S. Pat. No. 7,343,453, which is incorporated by reference herein.

Operations management can generally include monitoring and managing the health and performance of information management system **100** by, without limitation, performing error tracking, generating granular storage/performance metrics (e.g., job success/failure information, deduplication efficiency, etc.), generating storage modeling and costing information, and the like.

As an example, a storage manager **140** or other component in the information management system **100** may analyze traffic patterns and suggest or automatically route data via a particular route to e.g., certain facilitate storage and minimize congestion. In some embodiments, the system can generate predictions relating to storage operations or storage operation information. Such predictions described may be based on a trending analysis that may be used to predict various network operations or use of network resources such as network traffic levels, storage media use, use of bandwidth of communication links, use of media agent components, etc. Further examples of traffic analysis, trend analysis, prediction generation, and the like are described in U.S. Pat. No. 7,343,453, which is incorporated by reference herein.

In some configurations, a master storage manager **140** may track the status of a set of associated storage operation cells in a hierarchy of information management cells, such as the status of jobs, system components, system resources, and other items, by communicating with storage managers **140** (or other components) in the respective storage operation cells. Moreover, the master storage manager **140** may track the status of its associated storage operation cells and associated information management operations by receiving periodic status updates from the storage managers **140** (or other components) in the respective cells regarding jobs, system components, system resources, and other items. In some embodiments, a master storage manager **140** may store status information and other information regarding its associated storage operation cells and other system information in its index **150** (or other location).

The master storage manager **140** or other component in the system may also determine whether a storage-related criteria or other criteria is satisfied, and perform an action or trigger event (e.g., data migration) in response to the criteria being satisfied, such as where a storage threshold is met for a particular volume, or where inadequate protection exists for certain data. For instance, in some embodiments, the system uses data from one or more storage operation cells to advise users of risks or indicates actions that can be used to mitigate or otherwise minimize these risks, and in some embodiments, dynamically takes action to mitigate or minimize these risks. For example, an information management policy may specify certain requirements (e.g., that a storage device should maintain a certain amount of free space, that secondary copies should occur at a particular interval, that data should be aged and migrated to other storage after a particular period, that data on a secondary volume should always have a certain level of availability and be able to be restored within a given time period, that data on a secondary volume may be mirrored or otherwise migrated to a specified number of other volumes, etc.). If a risk condition or other criteria is triggered, the system can notify the user of

these conditions and may suggest (or automatically implement) an action to mitigate or otherwise address the condition or minimize risk. For example, the system may indicate that data from a primary copy **112** should be migrated to a secondary storage device **108** to free space on the primary storage device **104**. Examples of the use of risk factors and other triggering criteria are described in U.S. Pat. No. 7,343,453, which is incorporated by reference herein.

In some embodiments, the system **100** may also determine whether a metric or other indication satisfies a particular storage criteria and, if so, perform an action. For example, as previously described, a storage policy or other definition might indicate that a storage manager **140** should initiate a particular action if a storage metric or other indication drops below or otherwise fails to satisfy specified criteria such as a threshold of data protection. Examples of such metrics are described in U.S. Pat. No. 7,343,453, which is incorporated by reference herein.

In some embodiments, risk factors may be quantified into certain measurable service or risk levels for ease of comprehension. For example, certain applications and associated data may be considered to be more important by an enterprise than other data and services. Financial compliance data, for example, may be of greater importance than marketing materials, etc. Network administrators may assign priorities or “weights” to certain data or applications, corresponding to its importance (priority value). The level of compliance with the storage operations specified for these applications may also be assigned a certain value. Thus, the health, impact and overall importance of a service on an enterprise may be determined, for example, by measuring the compliance value and calculating the product of the priority value and the compliance value to determine the “service level” and comparing it to certain operational thresholds to determine if the operation is being performed within a specified data protection service level. Further examples of the service level determination are provided in U.S. Pat. No. 7,343,453, which is incorporated by reference herein.

The system **100** may additionally calculate data costing and data availability associated with information management operation cells according to an embodiment of the invention. For instance, data received from the cell may be used in conjunction with hardware-related information and other information about network elements to generate indications of costs associated with storage of particular data in the system or the availability of particular data in the system. In general, components in the system are identified and associated information is obtained (dynamically or manually). Characteristics or metrics associated with the network elements may be identified and associated with that component element for further use generating an indication of storage cost or data availability. Exemplary information generated could include how fast a particular department is using up available storage space, how long data would take to recover over a particular network pathway from a particular secondary storage device, costs over time, etc. Moreover, in some embodiments, such information may be used to determine or predict the overall cost associated with the storage of certain information. The cost associated with hosting a certain application may be based, at least in part, on the type of media on which the data resides. Storage devices may be assigned to a particular cost category which is indicative of the cost of storing information on that device. Further examples of costing techniques are described in U.S. Pat. No. 7,343,453, which is incorporated by reference herein.

Any of the above types of information (e.g., information related to trending, predictions, job, cell or component status, risk, service level, costing, etc.) can generally be provided to users via the user interface **158** in a single, integrated view or console. The console may support a reporting capability that allows for the generation of a variety of reports, which may be tailored to a particular aspect of information management. Report types may include: scheduling, event management, media management and data aging. Available reports may also include backup history, data aging history, auxiliary copy history, job history, library and drive, media in library, restore history, and storage policy. Such reports may be specified and created at a certain point in time as a network analysis, forecasting, or provisioning tool. Integrated reports may also be generated that illustrate storage and performance metrics, risks and storage costing information. Moreover, users may create their own reports based on specific needs.

The integrated user interface **158** can include an option to show a “virtual view” of the system that graphically depicts the various components in the system using appropriate icons. As one example, the user interface **158** may provide a graphical depiction of one or more primary storage devices **104**, the secondary storage devices **108**, data agents **142** and/or media agents **144**, and their relationship to one another in the information management system **100**. The operations management functionality can facilitate planning and decision-making. For example, in some embodiments, a user may view the status of some or all jobs as well as the status of each component of the information management system **100**. Users may then plan and make decisions based on this data. For instance, a user may view high-level information regarding storage operations for the information management system **100**, such as job status, component status, resource status (e.g., network pathways, etc.), and other information. The user may also drill down or use other means to obtain more detailed information regarding a particular component, job, or the like.

Further examples of some reporting techniques and associated interfaces providing an integrated view of an information management system are provided in U.S. Pat. No. 7,343,453, which is incorporated by reference herein.

The information management system **100** can also be configured to perform system-wide e-discovery operations in some embodiments. In general, e-discovery operations provide a unified collection and search capability for data in the system, such as data stored in the secondary storage devices **108** (e.g., backups, archives, or other secondary copies **116**). For example, the information management system **100** may construct and maintain a virtual repository for data stored in the information management system **100** that is integrated across source applications **110**, different storage device types, etc. According to some embodiments, e-discovery utilizes other techniques described herein, such as data classification and/or content indexing.

Information Management Policies

As indicated previously, an information management policy **148** can include a data structure or other information source that specifies a set of parameters (e.g., criteria and rules) associated with secondary copy or other information management operations.

One type of information management policy **148** is a storage policy. According to certain embodiments, a storage policy generally comprises a data structure or other information source that defines (or includes information sufficient to determine) a set of preferences or other criteria for performing information management operations. Storage

policies can include one or more of the following items: (1) what data will be associated with the storage policy; (2) a destination to which the data will be stored; (3) datapath information specifying how the data will be communicated to the destination; (4) the type of storage operation to be performed; and (5) retention information specifying how long the data will be retained at the destination.

As an illustrative example, data associated with a storage policy can be logically organized into groups. In some cases, these logical groupings can be referred to as “sub-clients”. A sub-client may represent static or dynamic associations of portions of a data volume. Sub-clients may represent mutually exclusive portions. Thus, in certain embodiments, a portion of data may be given a label and the association is stored as a static entity in an index, database or other storage location.

Sub-clients may also be used as an effective administrative scheme of organizing data according to data type, department within the enterprise, storage preferences, or the like. Depending on the configuration, sub-clients can correspond to files, folders, virtual machines, databases, etc. In one exemplary scenario, an administrator may find it preferable to separate e-mail data from financial data using two different sub-clients.

A storage policy can define where data is stored by specifying a target or destination storage device (or group of storage devices). For instance, where the secondary storage device **108** includes a group of disk libraries, the storage policy may specify a particular disk library for storing the sub-clients associated with the policy. As another example, where the secondary storage devices **108** include one or more tape libraries, the storage policy may specify a particular tape library for storing the sub-clients associated with the storage policy, and may also specify a drive pool and a tape pool defining a group of tape drives and a group of tapes, respectively, for use in storing the sub-client data. While information in the storage policy can be statically assigned in some cases, some or all of the information in the storage policy can also be dynamically determined based on criteria, which can be set forth in the storage policy. For instance, based on such criteria, a particular destination storage device(s) (or other parameter of the storage policy) may be determined based on characteristics associated with the data involved in a particular storage operation, device availability (e.g., availability of a secondary storage device **108** or a media agent **144**), network status and conditions (e.g., identified bottlenecks), user credentials, and the like).

Datapath information can also be included in the storage policy. For instance, the storage policy may specify network pathways and components to utilize when moving the data to the destination storage device(s). In some embodiments, the storage policy specifies one or more media agents **144** for conveying data (e.g., one or more sub-clients) associated with the storage policy between the source (e.g., one or more host client computing devices **102**) and destination (e.g., a particular target secondary storage device **108**).

A storage policy can also specify the type(s) of operations associated with the storage policy, such as a backup, archive, snapshot, auxiliary copy, or the like. Retention information can specify how long the data will be kept, depending on organizational needs (e.g., a number of days, months, years, etc.)

The information management policies **148** may also include one or more scheduling policies specifying when and how often to perform operations. Scheduling information may specify with what frequency (e.g., hourly, weekly, daily, event-based, etc.) or under what triggering conditions

secondary copy or other information management operations will take place. Scheduling policies in some cases are associated with particular components, such as particular logical groupings of data associated with a storage policy (e.g., a sub-client), client computing device **102**, and the like. In one configuration, a separate scheduling policy is maintained for particular logical groupings of data on a client computing device **102**. The scheduling policy specifies that those logical groupings are to be moved to secondary storage devices **108** every hour according to storage policies associated with the respective sub-clients.

When adding a new client computing device **102**, administrators can manually configure information management policies **148** and/or other settings, e.g., via the user interface **158**. However, this can be an involved process resulting in delays, and it may be desirable to begin data protecting operations quickly.

Thus, in some embodiments, the information management system **100** automatically applies a default configuration to client computing device **102**. As one example, when one or more data agent(s) **142** are installed on one or more client computing devices **102**, the installation script may register the client computing device **102** with the storage manager **140**, which in turn applies the default configuration to the new client computing device **102**. In this manner, data protection operations can begin substantially immediately. The default configuration can include a default storage policy, for example, and can specify any appropriate information sufficient to begin data protection operations. This can include a type of data protection operation, scheduling information, a target secondary storage device **108**, data path information (e.g., a particular media agent **144**), and the like.

Other types of information management policies **148** are possible. For instance, the information management policies **148** can also include one or more audit or security policies. An audit policy is a set of preferences, rules and/or criteria that protect sensitive data in the information management system **100**. For example, an audit policy may define “sensitive objects” as files or objects that contain particular keywords (e.g., “confidential,” or “privileged”) and/or are associated with particular keywords (e.g., in metadata) or particular flags (e.g., in metadata identifying a document or email as personal, confidential, etc.).

An audit policy may further specify rules for handling sensitive objects. As an example, an audit policy may require that a reviewer approve the transfer of any sensitive objects to a cloud storage site, and that if approval is denied for a particular sensitive object, the sensitive object should be transferred to a local primary storage device **104** instead. To facilitate this approval, the audit policy may further specify how a secondary storage computing device **106** or other system component should notify a reviewer that a sensitive object is slated for transfer.

In some implementations, the information management policies **148** may include one or more provisioning policies. A provisioning policy can include a set of preferences, priorities, rules, and/or criteria that specify how client computing devices **102** (or groups thereof) may utilize system resources, such as available storage on cloud storage and/or network bandwidth. A provisioning policy specifies, for example, data quotas for particular client computing devices **102** (e.g., a number of gigabytes that can be stored monthly, quarterly or annually). The storage manager **140** or other components may enforce the provisioning policy. For instance, the media agents **144** may enforce the policy when transferring data to secondary storage devices **108**. If a client computing device **102** exceeds a quota, a budget for the

client computing device **102** (or associated department) is adjusted accordingly or an alert may trigger.

While the above types of information management policies **148** have been described as separate policies, one or more of these can be generally combined into a single information management policy **148**. For instance, a storage policy may also include or otherwise be associated with one or more scheduling, audit, or provisioning policies. Moreover, while storage policies are typically associated with moving and storing data, other policies may be associated with other types of information management operations. The following is a non-exhaustive list of items the information management policies **148** may specify:

schedules or other timing information, e.g., specifying when and/or how often to perform information management operations;

the type of copy **116** (e.g., type of secondary copy) and/or copy format (e.g., snapshot, backup, archive, HSM, etc.);

a location or a class or quality of storage for storing secondary copies **116** (e.g., one or more particular secondary storage devices **108**);

preferences regarding whether and how to encrypt, compress, deduplicate, or otherwise modify or transform secondary copies **116**;

which system components and/or network pathways (e.g., preferred media agents **144**) should be used to perform secondary storage operations;

resource allocation between different computing devices or other system components used in performing information management operations (e.g., bandwidth allocation, available storage capacity, etc.);

whether and how to synchronize or otherwise distribute files or other data objects across multiple computing devices or hosted services; and

retention information specifying the length of time primary data **112** and/or secondary copies **116** should be retained, e.g., in a particular class or tier of storage devices, or within the information management system **100**.

Policies can additionally specify or depend on a variety of historical or current criteria that may be used to determine which rules to apply to a particular data object, system component, or information management operation, such as:

frequency with which primary data **112** or a secondary copy **116** of a data object or metadata has been or is predicted to be used, accessed, or modified;

time-related factors (e.g., aging information such as time since the creation or modification of a data object);

deduplication information (e.g., hashes, data blocks, deduplication block size, deduplication efficiency or other metrics);

an estimated or historic usage or cost associated with different components (e.g., with secondary storage devices **108**);

the identity of users, applications **110**, client computing devices **102** and/or other computing devices that created, accessed, modified, or otherwise utilized primary data **112** or secondary copies **116**;

a relative sensitivity (e.g., confidentiality) of a data object, e.g., as determined by its content and/or metadata;

the current or historical storage capacity of various storage devices;

the current or historical network capacity of network pathways connecting various components within the storage operation cell;

access control lists or other security information; and the content of a particular data object (e.g., its textual content) or of metadata associated with the data object.

Exemplary Storage Policy and Secondary Storage Operations

FIG. 1E shows a data flow data diagram depicting performance of storage operations by an embodiment of an information management system **100**, according to an exemplary storage policy **148A**. The information management system **100** includes a storage manager **140**, a client computing device **102** having a file system data agent **142A** and an email data agent **142B** residing thereon, a primary storage device **104**, two media agents **144A**, **144B**, and two secondary storage devices **108A**, **108B**: a disk library **108A** and a tape library **108B**. As shown, the primary storage device **104** includes primary data **112A**, **112B** associated with a logical grouping of data associated with a file system) and a logical grouping of data associated with email data, respectively. Although for simplicity the logical grouping of data associated with the file system is referred to as a file system sub-client, and the logical grouping of data associated with the email data is referred to as an email sub-client, the techniques described with respect to FIG. 1E can be utilized in conjunction with data that is organized in a variety of other manners.

As indicated by the dashed box, the second media agent **144B** and the tape library **108B** are “off-site”, and may therefore be remotely located from the other components in the information management system **100** (e.g., in a different city, office building, etc.). Indeed, “off-site” may refer to a magnetic tape located in storage, which must be manually retrieved and loaded into a tape drive to be read. In this manner, information stored on the tape library **108B** may provide protection in the event of a disaster or other failure.

The file system sub-client and its associated primary data **112A** in certain embodiments generally comprise information generated by the file system and/or operating system of the client computing device **102**, and can include, for example, file system data (e.g., regular files, file tables, mount points, etc.), operating system data (e.g., registries, event logs, etc.), and the like. The e-mail sub-client, on the other hand, and its associated primary data **112B**, include data generated by an e-mail client application operating on the client computing device **102**, and can include mailbox information, folder information, emails, attachments, associated database information, and the like. As described above, the sub-clients can be logical containers, and the data included in the corresponding primary data **112A**, **112B** may or may not be stored contiguously.

The exemplary storage policy **148A** includes backup copy preferences or rule set **160**, disaster recovery copy preferences rule set **162**, and compliance copy preferences or rule set **164**. The backup copy rule set **160** specifies that it is associated with a file system sub-client **166** and an email sub-client **168**. Each of these sub-clients **166**, **168** are associated with the particular client computing device **102**. The backup copy rule set **160** further specifies that the backup operation will be written to the disk library **108A**, and designates a particular media agent **144A** to convey the data to the disk library **108A**. Finally, the backup copy rule set **160** specifies that backup copies created according to the rule set **160** are scheduled to be generated on an hourly basis and to be retained for 30 days. In some other embodiments, scheduling information is not included in the storage policy **148A**, and is instead specified by a separate scheduling policy.

The disaster recovery copy rule set **162** is associated with the same two sub-clients **166**, **168**. However, the disaster

recovery copy rule set **162** is associated with the tape library **1088**, unlike the backup copy rule set **160**. Moreover, the disaster recovery copy rule set **162** specifies that a different media agent **144B** than the media agent **144A** associated with the backup copy rule set **160** will be used to convey the data to the tape library **1088**. As indicated, disaster recovery copies created according to the rule set **162** will be retained for 60 days, and will be generated on a daily basis. Disaster recovery copies generated according to the disaster recovery copy rule set **162** can provide protection in the event of a disaster or other data-loss event that would affect the backup copy **116A** maintained on the disk library **108A**.

The compliance copy rule set **164** is only associated with the email sub-client **168**, and not the file system sub-client **166**. Compliance copies generated according to the compliance copy rule set **164** will therefore not include primary data **112A** from the file system sub-client **166**. For instance, the organization may be under an obligation to store and maintain copies of email data for a particular period of time (e.g., 10 years) to comply with state or federal regulations, while similar regulations do not apply to the file system data. The compliance copy rule set **164** is associated with the same tape library **1088** and media agent **1448** as the disaster recovery copy rule set **162**, although a different storage device or media agent could be used in other embodiments. Finally, the compliance copy rule set **164** specifies that copies generated under the compliance copy rule set **164** will be retained for 10 years, and will be generated on a quarterly basis.

At step 1, the storage manager **140** initiates a backup operation according to the backup copy rule set **160**. For instance, a scheduling service running on the storage manager **140** accesses scheduling information from the backup copy rule set **160** or a separate scheduling policy associated with the client computing device **102**, and initiates a backup copy operation on an hourly basis. Thus, at the scheduled time slot the storage manager **140** sends instructions to the client computing device **102** to begin the backup operation.

At step 2, the file system data agent **142A** and the email data agent **142B** residing on the client computing device **102** respond to the instructions received from the storage manager **140** by accessing and processing the primary data **112A**, **112B** involved in the copy operation from the primary storage device **104**. Because the operation is a backup copy operation, the data agent(s) **142A**, **142B** may format the data into a backup format or otherwise process the data.

At step 3, the client computing device **102** communicates the retrieved, processed data to the first media agent **144A**, as directed by the storage manager **140**, according to the backup copy rule set **160**. In some other embodiments, the information management system **100** may implement a load-balancing, availability-based, or other appropriate algorithm to select from the available set of media agents **144A**, **144B**. Regardless of the manner the media agent **144A** is selected, the storage manager **140** may further keep a record in the storage manager database **146** of the association between the selected media agent **144A** and the client computing device **102** and/or between the selected media agent **144A** and the backup copy **116A**.

The target media agent **144A** receives the data from the client computing device **102**, and at step 4 conveys the data to the disk library **108A** to create the backup copy **116A**, again at the direction of the storage manager **140** and according to the backup copy rule set **160**. The secondary storage device **108A** can be selected in other ways. For instance, the media agent **144A** may have a dedicated association with a particular secondary storage device(s), or

the storage manager **140** or media agent **144A** may select from a plurality of secondary storage devices, e.g., according to availability, using one of the techniques described in U.S. Pat. No. 7,246,207, which is incorporated by reference herein.

The media agent **144A** can also update its index **153** to include data and/or metadata related to the backup copy **116A**, such as information indicating where the backup copy **116A** resides on the disk library **108A**, data and metadata for cache retrieval, etc. After the 30 day retention period expires, the storage manager **140** instructs the media agent **144A** to delete the backup copy **116A** from the disk library **108A**. The storage manager **140** may similarly update its index **150** to include information relating to the storage operation, such as information relating to the type of storage operation, a physical location associated with one or more copies created by the storage operation, the time the storage operation was performed, status information relating to the storage operation, the components involved in the storage operation, and the like. In some cases, the storage manager **140** may update its index **150** to include some or all of the information stored in the index **153** of the media agent **144A**.

At step 5, the storage manager **140** initiates the creation of a disaster recovery copy **1168** according to the disaster recovery copy rule set **162**. For instance, at step 6, based on instructions received from the storage manager **140** at step 5, the specified media agent **1448** retrieves the most recent backup copy **116A** from the disk library **108A**.

At step 7, again at the direction of the storage manager **140** and as specified in the disaster recovery copy rule set **162**, the media agent **144B** uses the retrieved data to create a disaster recovery copy **1168** on the tape library **1088**. In some cases, the disaster recovery copy **1168** is a direct, mirror copy of the backup copy **116A**, and remains in the backup format. In other embodiments, the disaster recovery copy **1168** may be generated in some other manner, such as by using the primary data **112A**, **1128** from the primary storage device **104** as source data. The disaster recovery copy operation is initiated once a day and the disaster recovery copies **1168** are deleted after 60 days.

At step 8, the storage manager **140** initiates the creation of a compliance copy **116C**, according to the compliance copy rule set **164**. For instance, the storage manager **140** instructs the media agent **144B** to create the compliance copy **116C** on the tape library **1088** at step 9, as specified in the compliance copy rule set **164**. In the example, the compliance copy **116C** is generated using the disaster recovery copy **1168**. In other embodiments, the compliance copy **116C** is instead generated using either the primary data **1128** corresponding to the email sub-client or using the backup copy **116A** from the disk library **108A** as source data. As specified, in the illustrated example, compliance copies **116C** are created quarterly, and are deleted after ten years.

While not shown in FIG. 1E, at some later point in time, a restore operation can be initiated involving one or more of the secondary copies **116A**, **1168**, **116C**. As one example, a user may manually initiate a restore of the backup copy **116A** by interacting with the user interface **158** of the storage manager **140**. The storage manager **140** then accesses data in its index **150** (and/or the respective storage policy **148A**) associated with the selected backup copy **116A** to identify the appropriate media agent **144A** and/or secondary storage device **108A**.

In other cases, a media agent may be selected for use in the restore operation based on a load balancing algorithm, an availability based algorithm, or other criteria. The selected

media agent **144A** retrieves the data from the disk library **108A**. For instance, the media agent **144A** may access its index **153** to identify a location of the backup copy **116A** on the disk library **108A**, or may access location information residing on the disk **108A** itself.

When the backup copy **116A** was recently created or accessed, the media agent **144A** accesses a cached version of the backup copy **116A** residing in the index **153**, without having to access the disk library **108A** for some or all of the data. Once it has retrieved the backup copy **116A**, the media agent **144A** communicates the data to the source client computing device **102**. Upon receipt, the file system data agent **142A** and the email data agent **142B** may unpackage (e.g., restore from a backup format to the native application format) the data in the backup copy **116A** and restore the unpackaged data to the primary storage device **104**.

Exemplary Applications of Storage Policies

The storage manager **140** may permit a user to specify aspects of the storage policy **148A**. For example, the storage policy can be modified to include information governance policies to define how data should be managed in order to comply with a certain regulation or business objective. The various policies may be stored, for example, in the database **146**. An information governance policy may comprise a classification policy, which is described herein. An information governance policy may align with one or more compliance tasks that are imposed by regulations or business requirements. Examples of information governance policies might include a Sarbanes-Oxley policy, a HIPAA policy, an electronic discovery (E-Discovery) policy, and so on.

Information governance policies allow administrators to obtain different perspectives on all of an organization's online and offline data, without the need for a dedicated data silo created solely for each different viewpoint. As described previously, the data storage systems herein build a centralized index that reflects the contents of a distributed data set that spans numerous clients and storage devices, including both primary and secondary copies, and online and offline copies. An organization may apply multiple information governance policies in a top-down manner over that unified data set and indexing schema in order to permit an organization to view and manipulate the single data set through different lenses, each of which is adapted to a particular compliance or business goal. Thus, for example, by applying an E-discovery policy and a Sarbanes-Oxley policy, two different groups of users in an organization can conduct two very different analyses of the same underlying physical set of data copies, which may be distributed throughout the organization.

A classification policy defines a taxonomy of classification terms or tags relevant to a compliance task and/or business objective. A classification policy may also associate a defined tag with a classification rule. A classification rule defines a particular combination of data criteria, such as users who have created, accessed or modified a document or data object; file or application types; content or metadata keywords; clients or storage locations; dates of data creation and/or access; review status or other status within a workflow (e.g., reviewed or un-reviewed); modification times or types of modifications; and/or any other data attributes. A classification rule may also be defined using other classification tags in the taxonomy. The various criteria used to define a classification rule may be combined in any suitable fashion, for example, via Boolean operators, to define a complex classification rule. As an example, an E-discovery classification policy might define a classification tag "privileged" that is associated with documents or data objects that

(1) were created or modified by legal department staff, (2) were sent to or received from outside counsel via email, and/or (3) contain one of the following keywords: "privileged" or "attorney," "counsel", or other terms.

One specific type of classification tag, which may be added to an index at the time of indexing, is an entity tag. An entity tag may be, for example, any content that matches a defined data mask format. Examples of entity tags might include, e.g., social security numbers (e.g., any numerical content matching the formatting mask XXX-XX-XXXX), credit card numbers (e.g., content having a 13-16 digit string of numbers), SKU numbers, product numbers, etc.

A user may define a classification policy by indicating criteria, parameters or descriptors of the policy via a graphical user interface that provides facilities to present information and receive input data, such as a form or page with fields to be filled in, pull-down menus or entries allowing one or more of several options to be selected, buttons, sliders, hypertext links or other known user interface tools for receiving user input. For example, a user may define certain entity tags, such as a particular product number or project ID code that is relevant in the organization.

In some implementations, the classification policy can be implemented using cloud-based techniques. For example, the storage devices may be cloud storage devices, and the storage manager **140** may execute cloud service provider API over a network to classify data stored on cloud storage devices.

Exemplary Secondary Copy Formatting

The formatting and structure of secondary copies **116** can vary, depending on the embodiment. In some cases, secondary copies **116** are formatted as a series of logical data units or "chunks" (e.g., 512 MB, 1 GB, 2 GB, 4 GB, or 8 GB chunks). This can facilitate efficient communication and writing to secondary storage devices **108**, e.g., according to resource availability. For example, a single secondary copy **116** may be written on a chunk-by-chunk basis to a single secondary storage device **108** or across multiple secondary storage devices **108**. In some cases, users can select different chunk sizes, e.g., to improve throughput to tape storage devices.

Generally, each chunk can include a header and a payload. The payload can include files (or other data units) or subsets thereof included in the chunk, whereas the chunk header generally includes metadata relating to the chunk, some or all of which may be derived from the payload. For example, during a secondary copy operation, the media agent **144**, storage manager **140**, or other component may divide the associated files into chunks and generate headers for each chunk by processing the constituent files.

The headers can include a variety of information such as file identifier(s), volume(s), offset(s), or other information associated with the payload data items, a chunk sequence number, etc. Importantly, in addition to being stored with the secondary copy **116** on the secondary storage device **108**, the chunk headers can also be stored to the index **153** of the associated media agent(s) **144** and/or the index **150**. This is useful in some cases for providing faster processing of secondary copies **116** during restores or other operations. In some cases, once a chunk is successfully transferred to a secondary storage device **108**, the secondary storage device **108** returns an indication of receipt, e.g., to the media agent **144** and/or storage manager **140**, which may update their respective indexes **153**, **150** accordingly. During restore, chunks may be processed (e.g., by the media agent **144**) according to the information in the chunk header to reassemble the files.

Data can also be communicated within the information management system **100** in data channels that connect the client computing devices **102** to the secondary storage devices **108**. These data channels can be referred to as “data streams”, and multiple data streams can be employed to parallelize an information management operation, improving data transfer rate, among providing other advantages. Example data formatting techniques including techniques involving data streaming, chunking, and the use of other data structures in creating copies (e.g., secondary copies) are described in U.S. Pat. Nos. 7,315,923 and 8,156,086, and 8,578,120, each of which is incorporated by reference herein.

FIGS. 1F and 1G are diagrams of example data streams **170** and **171**, respectively, which may be employed for performing data storage operations. Referring to FIG. 1F, the data agent **142** forms the data stream **170** from the data associated with a client computing device **102** (e.g., primary data **112**). The data stream **170** is composed of multiple pairs of stream header **172** and stream data (or stream payload) **174**. The data streams **170** and **171** shown in the illustrated example are for a single-instanced storage operation, and a stream payload **174** therefore may include both single-instance (“SI”) data and/or non-SI data. A stream header **172** includes metadata about the stream payload **174**. This metadata may include, for example, a length of the stream payload **174**, an indication of whether the stream payload **174** is encrypted, an indication of whether the stream payload **174** is compressed, an archive file identifier (ID), an indication of whether the stream payload **174** is single instanceable, and an indication of whether the stream payload **174** is a start of a block of data.

Referring to FIG. 1G, the data stream **171** has the stream header **172** and stream payload **174** aligned into multiple data blocks. In this example, the data blocks are of size 64 KB. The first two stream header **172** and stream payload **174** pairs comprise a first data block of size 64 KB. The first stream header **172** indicates that the length of the succeeding stream payload **174** is 63 KB and that it is the start of a data block. The next stream header **172** indicates that the succeeding stream payload **174** has a length of 1 KB and that it is not the start of a new data block. Immediately following stream payload **174** is a pair comprising an identifier header **176** and identifier data **178**. The identifier header **176** includes an indication that the succeeding identifier data **178** includes the identifier for the immediately previous data block. The identifier data **178** includes the identifier that the data agent **142** generated for the data block. The data stream **171** also includes other stream header **172** and stream payload **174** pairs, which may be for SI data and/or for non-SI data.

FIG. 1H is a diagram illustrating the data structures **180** that may be used to store blocks of SI data and non-SI data on the storage device (e.g., secondary storage device **108**). According to certain embodiments, the data structures **180** do not form part of a native file system of the storage device. The data structures **180** include one or more volume folders **182**, one or more chunk folders **184/185** within the volume folder **182**, and multiple files within the chunk folder **184**. Each chunk folder **184/185** includes a metadata file **186/187**, a metadata index file **188/189**, one or more container files **190/191/193**, and a container index file **192/194**. The metadata file **186/187** stores non-SI data blocks as well as links to SI data blocks stored in container files. The metadata index file **188/189** stores an index to the data in the metadata file **186/187**. The container files **190/191/193** store SI data blocks. The container index file **192/194** stores an index to

the container files **190/191/193**. Among other things, the container index file **192/194** stores an indication of whether a corresponding block in a container file **190/191/193** is referred to by a link in a metadata file **186/187**. For example, data block B2 in the container file **190** is referred to by a link in the metadata file **187** in the chunk folder **185**. Accordingly, the corresponding index entry in the container index file **192** indicates that the data block B2 in the container file **190** is referred to. As another example, data block B1 in the container file **191** is referred to by a link in the metadata file **187**, and so the corresponding index entry in the container index file **192** indicates that this data block is referred to.

As an example, the data structures **180** illustrated in FIG. 1H may have been created as a result of two storage operations involving two client computing devices **102**. For example, a first storage operation on a first client computing device **102** could result in the creation of the first chunk folder **184**, and a second storage operation on a second client computing device **102** could result in the creation of the second chunk folder **185**. The container files **190/191** in the first chunk folder **184** would contain the blocks of SI data of the first client computing device **102**. If the two client computing devices **102** have substantially similar data, the second storage operation on the data of the second client computing device **102** would result in the media agent **144** storing primarily links to the data blocks of the first client computing device **102** that are already stored in the container files **190/191**. Accordingly, while a first storage operation may result in storing nearly all of the data subject to the storage operation, subsequent storage operations involving similar data may result in substantial data storage space savings, because links to already stored data blocks can be stored instead of additional instances of data blocks.

If the operating system of the secondary storage computing device **106** on which the media agent **144** resides supports sparse files, then when the media agent **144** creates container files **190/191/193**, it can create them as sparse files. As previously described, a sparse file is type of file that may include empty space (e.g., a sparse file may have real data within it, such as at the beginning of the file and/or at the end of the file, but may also have empty space in it that is not storing actual data, such as a contiguous range of bytes all having a value of zero). Having the container files **190/191/193** be sparse files allows the media agent **144** to free up space in the container files **190/191/193** when blocks of data in the container files **190/191/193** no longer need to be stored on the storage devices. In some examples, the media agent **144** creates a new container file **190/191/193** when a container file **190/191/193** either includes 100 blocks of data or when the size of the container file **190** exceeds 50 MB. In other examples, the media agent **144** creates a new container file **190/191/193** when a container file **190/191/193** satisfies other criteria (e.g., it contains from approximately 100 to approximately 1000 blocks or when its size exceeds approximately 50 MB to 1 GB).

In some cases, a file on which a storage operation is performed may comprise a large number of data blocks. For example, a 100 MB file may be comprised in 400 data blocks of size 256 KB. If such a file is to be stored, its data blocks may span more than one container file, or even more than one chunk folder. As another example, a database file of 20 GB may comprise over 40,000 data blocks of size 512 KB. If such a database file is to be stored, its data blocks will likely span multiple container files, multiple chunk folders, and potentially multiple volume folders. As described in detail herein, restoring such files may thus requiring access-

ing multiple container files, chunk folders, and/or volume folders to obtain the requisite data blocks.

Example Scalable Information Management System

FIG. 2 is a block diagram illustrating an example of a scalable information management system 200. In some embodiments, the scalable information management system 200 can include some or all of the embodiments described previously with respect to the information management systems 100. To simplify discussion, reference numbers are re-used to indicate correspondence between certain referenced elements of the information management system 200 and the information management system 100. Further, to simplify the illustration, certain elements are omitted from the illustration of the information management system 200, but which may be included in certain embodiments. For example, although the storage manager 140 of the information management system 200 omits a job agent, the storage manager 140 as illustrated in FIG. 1C may include a jobs agent 156.

The information management system 200 can implement a number of processes in addition to those that have been previously described. For example, the information management system 200 may implement a process to reduce the proliferation of virtual machines (VMs) by allocating job requests among existing virtual machines. As a second example, the information management system 200 can implement a process for grouping virtual machines and/or virtual machine host systems or provider systems. As a third example, the information management system 200 can implement a load-balancing process or distribute work among virtual server agents. Although not limited as such, this load-balancing process may be used to distribute the load during a backup process. These example processes are described in more detail below with reference to FIGS. 3-5. Further, the example processes are described with respect to a number of example systems of the information management system 200. However, other systems may implement the processes including some of the systems that have previously been described. Some of the example systems will now be described.

As previously described, the information management system 200 may include a number of client computing devices 102. In addition to the applications 110 and the data agents 142 (not shown in FIG. 2), the computing devices may include a number of virtual machines 204. The virtual machines 204 may include any type of virtual machine created for any purpose. For example, the virtual machines 204 may be Windows-based virtual machines, UNIX-based virtual machines, or Apple OS-based virtual machines. Further, the virtual machines 204 may be for facilitating application access, emulating hardware, accessing or managing different databases, performing data management, sharing computing resources among multiple user, or for any other purpose.

In addition to the virtual machines 204, the client computing devices 102 may include a virtual machine monitor 206 or a hypervisor. The virtual machine monitor 206 may be any type of virtual machine monitor for managing the virtual machines 204. For example, the virtual machine monitor 206 may be a native or bare metal hypervisor (e.g., Oracle VM Server or Microsoft Hyper-V) or a hosted hypervisor run within an operating system environment (e.g., VMware Workstation or VirtualBox). In some cases, multiple virtual machine monitors 206 may be included by the client computing devices 102. For example, a virtual machine monitor 206 may be included for one type or group

of virtual machines and another virtual machine monitor 206 may be included for another type or group of virtual machines.

In addition to the client computing devices 102, the information management system 200 includes a number of server computing devices 212. Like the client computing devices 102, a server computing device 212 may include a number of virtual machines 204 and a virtual machine monitor 206. Although the terms client and server are used with respect to the computing devices of the primary storage subsystem 117, this disclosure is not limited to a client/server computing infrastructure, but may include other types of computing systems and/or network configurations.

As illustrated in FIG. 2, the information management system 200 may include a number of virtual server agents 210A, 210B, 210C (referred to collectively as virtual server agents or VSAs 210). The virtual server agents 210 can include any type of agent or system for managing operations between the primary storage subsystem 117 and the secondary storage subsystem 118 with respect to the virtual machines 204. For example, the virtual server agents 210 may facilitate or implement a backup and/or restoration process for backing/restoring files of a virtual machine 204 or a virtual disk to/from the secondary storage subsystem 118. As another example, the virtual server agents 210 may facilitate deduplication or encryption/decryption of files or virtual disks to be stored to or restored from the secondary storage subsystem 118. The virtual server agents may be standalone computing systems (e.g., the virtual server agent 210B). Alternatively, the virtual server agents may be a hardware and/or software module included with one or more of the computing systems hosting the virtual machines (e.g., the virtual server agents 210A of the server computing devices 212).

In some embodiments, the virtual server agents 210 may be grouped. For example, virtual server agents 210A may be grouped and may be configured to manage operations and interactions between the server computing devices 212 and the secondary storage subsystem 118. As a second example, the virtual server agents 210C may be grouped and may be configured to manage operations and interactions between the server computing device 212 and the secondary storage subsystem 118 instead of or in addition to the virtual server agents 210A.

Each of the virtual server agents 210 may include one or more data connections, data paths, data streams, or streams to the media agents 144 hosted by the secondary storage computing devices 106. Via the streams, the virtual machines, files accessed via the virtual machines, and/or virtual disks of the virtual machines may be communicated to/from the secondary storage subsystem 118. These data connections or communication paths are illustrated by the solid line arrows between the virtual server agents 210 and the secondary storage computing devices 106 indicating paths for data transfer. Although not illustrated, in some cases, commands may also be communicated from a system of the secondary storage subsystem 118 to a virtual server agent 210 via a command communication channel. In some cases, each communication channel between a virtual server agent 210 and a media agent 144 may include a single data stream. Alternatively, a communication channel may include multiple data streams.

Information management system 200 may also include one or more VSA coordinators 202A and 202B (referred to collectively as VSA coordinators 202). The VSA coordinators 202 may manage or help distribute job requests among virtual server agents 210 assigned to one or more virtual

machines **204**. For example, the storage manager **140** may determine, based for example on a backup policy stored at the repository **146**, that the virtual machines **204** of the server computing devices **212** are scheduled for backup. Consequently, the storage manager **140** may provide a job request to the VSA coordinator **202A** to backup the virtual machines **204** of the server computing devices **212**. The VSA coordinator **202A** may select one or more VSAs (e.g., the virtual server agents **210C**) assigned to the virtual machines **204** of the server computing devices **212** to facilitate the scheduled backup to the secondary storage subsystem **118**. Further, in some cases, the VSA coordinator **202A** may distribute the virtual machines **204** of the server computing devices **212** among the virtual server agents included in the group of virtual server agents **210C**.

As indicated by the dashed line arrows between the virtual server agents **210** and the VSA coordinators **202**, a virtual server agent **210** may receive commands and/or control data from the VSA coordinators **202**. In some embodiments, the VSA coordinators **202** may be standalone systems as indicated by the VSA coordinator **202B**. In other embodiments, the VSA coordinators **202** may be included as part of the storage manager **140** as illustrated by the VSA coordinator **202A**.

In some instances, a virtual server agent **210** may serve as a VSA coordinator **202**. For example, when a job request is received by a group of virtual server agents associated with a set of virtual machines, a virtual server agent from the group of virtual server agents may designate itself or may be designated as the VSA coordinator using a leader selection algorithm. This leader selection algorithm may include any type of algorithm for selecting a leader among a group of systems. For example, the leader selection algorithm may be based on a round robin algorithm where the leader rotates for each new job request. As another example, the leader selection algorithm may determine the least busy virtual server agent and select that server agent as the VSA coordinator for the current job request, or for a particular time period or set of job requests.

As previously stated, virtual machines **204** and/or virtual machine provider system or hosts (e.g., the server computing devices **212**) may be categorized or grouped together. Further, a set of virtual server agents **210** may be grouped together based on characteristics or the virtual server agents. In some cases, the set of virtual server agents may be grouped solely based on their assignment to the same set of virtual machines and/or virtual machine provider systems or hosts.

In some embodiments, the VSA coordinator **202** may group the virtual machine **204** and/or the virtual machine provider systems (e.g., the server computing devices **212**). For example, the VSA coordinator **202** may identify a number of characteristics associated with one or more computing devices configured to host virtual machines **204**. The VSA coordinator **202** may then group the computing devices based on the identified characteristics. After grouping the computing devices, the VSA coordinator may assign one or more virtual server agents to the group of computing devices based on characteristics of the virtual server agents (e.g., geographic location of the virtual server agents, and network location of the virtual server agents, or the media agents to which the virtual server agents have established communication streams). In some embodiments, virtual machines **204** and/or virtual machine provider systems may be grouped by a virtual server agent, which in some cases may be selected by the VSA coordinator **202**.

Further, in some embodiments the VSA coordinator **202** may perform a load-balancing operation with respect to the virtual server agents **210**. For example, when an operation, such as a backup operation, is initiated, the VSA coordinator **202** may distribute the virtual machines **204** among the assigned virtual server agents **210** based on the capacity of each of the virtual server agents **210**, the size of each of the virtual machines **204**, and/or the number of streams between each of the virtual server agents and the media agents of the secondary storage subsystem **118**.

As illustrated in FIG. 2, the information management system **200** may further include a virtual machine management interface **245**, which when implemented in software may sometimes be referred to as VM navigation software. The VM management interface **245** can include any system implemented in hardware and/or software that can provide a centralized platform for managing virtual infrastructure. The centralized platform allows visibility into the configuration of the virtual machines **204** within the information management system **200**. The VM management interface **245** can store information about the structural relationship between physical servers or hosts, resource pools or store devices (e.g., the primary storage devices **104**), and virtual machines **204** in the information management system **200**. Examples of a VM management interface **245** include VMware vCenter™, Microsoft System Center Virtual Machine Manager®, and the like. The VM management interface **245** can interface with the VM monitors **206** to retrieve information about the virtual machines **204**, which may be stored in a database or repository associated with the VM management interface **245**. Examples of the stored information can include a name, an address, other identifying information of each host and primary storage device **104** associated with the virtual machines **204** in the information management system **200**, status information for each virtual machine **204**, type of each virtual machine **204**, identity of associated virtual server agents **210**, an assigned secondary storage computing device **106**, if any, and any other information that may be associated with the virtual machines **204**.

In some embodiments, the VM management interface **245** may present information associated with the VMs to a user (e.g., an administrator). Further, the VM management interface **245** may be utilized by a user to manage the virtual machines **204** across one or more host systems (e.g., the client computing devices **102** and/or server computing devices **212**).

Further, in some implementations, the VM management interface **245** may include the VSA coordinator **202**. Alternatively, the VM management interface **245** may be separate from the VSA coordinator **202**. In some such cases, the VM management interface **245** may interact with the VSA coordinator **202** to obtain information regarding relating to the VSA groups and/or VSAs **210**. Further, in some cases, the VM management interface **245** may be used to configure, manage, or control the VSAs **210** by interacting with the VSA coordinators **202**. In some embodiments, one of the virtual machine management interface **245** and the VSA coordinator **202** may be optional.

To simplify the illustration, the VM management interface **245** is illustrated in FIG. 2 as communicating with the storage manager **140**, the client computing devices **102**, and the virtual server agent **210B**. However, the VM management interface **245** is not limited as such. Instead, as described above, the VM management interface **245** may communicate with one or more of the VSA coordinators **202**, virtual server agents **210**, and/or virtual machine monitors **206**.

Example Virtual Machine Job Allocation Process

FIG. 3 illustrates a flowchart for an example virtual machine job allocation process 300. The process 300 can be implemented, at least in part, by any system that can allocate a job request to a virtual machine from among a set of virtual machines based at least partially on the load of each of the virtual machines from the set of virtual machines. Advantageously, in certain embodiments, by allocating job requests to existing virtual machines, the proliferation of new virtual machines may be reduced in the information management system 200. By reducing the proliferation of new virtual machines, resource consumption can consequently be reduced. For example, in some cases, the number of server computing devices 212 maintained in the information management system 200 may be reduced. In certain embodiments, the process 300 may be implemented, in whole or in part, by a storage manager 140, a VSA coordinator 202, a VM management interface 245, and/or a virtual server agent 210, to name a few. Although any number of systems, in whole or in part, can implement the process 300, to simplify discussion, portions of the process 300 will be described with reference to particular systems.

The process 300 begins at block 302 where, for example, a VSA (e.g., VSA 210B) receives a new job request. This job request may be received from a VSA coordinator (e.g., VSA coordinator 202A), from a jobs agent 156, from a VM management interface 245, or from any other system that may provide a job request. In some embodiments, the job request may be received from another VSA, such as a VSA serving as a VSA coordinator for a group of virtual machines. In some cases, the job request may be received by the VSA from itself, such as when the VSA serves as a VSA coordinator and in its capacity as VSA coordinator selects itself to process the job request. The VSA receiving the job request may be one of a number of VSAs assigned to at least partially manage an assigned set of virtual machines.

At block 304, the VSA determines a load for the virtual machines 204 assigned to the VSA. The VSA may determine the load of the virtual machines by, for example, accessing one or more virtual machine monitors 206 associated with the virtual machines 204 and/or one or more data agents associated with the virtual machine. Alternatively, or in addition, the VSA may determine the load for one or more of the virtual machines 204 by tracking and/or accessing a table that tracks jobs or processes that have been assigned to the one or more virtual machines 204. In certain embodiments, the VSA is able to track the jobs assigned to the virtual machines 204 because the VSA assigns to and/or filters job requests for the virtual machines 204. In certain embodiments, the VSA may determine the load for a subset of virtual machines that are assigned to the VSA, which may be assigned based on a type of the virtual machine 204 and/or the type of jobs to be allocated to the virtual machines 204. In other words, in some cases, a subset of virtual machines may be allocated for one particular job type and another subset of virtual machines may be allocated for another job type. In such cases, the VSA may limit its analysis of load to the virtual machines that are allocated to the job type of the job request received at the block 302. In some embodiments, the VSA may determine the load for the virtual machines 204 by accessing or requesting load or status information from the VM management interface 245.

The VSA determines at decision block 306 whether there are any virtual machines with a load that satisfies a threshold. Further, in some cases, the decision block 306 may determine whether the VMs with a load that satisfies the threshold is a non-empty set. In some cases, determining

whether the load satisfies a threshold may include determining whether the load is below a threshold and/or whether the load is at or below a threshold. In some instances, the load may be preset by, for example, an administrator of the information management system 200. Further, in some cases, the threshold may be based at least partially on the job request or a type of job associated with the job request. Moreover, in some cases, the threshold may be based at least partially on the anticipated job load from the job request received at the block 302.

If the VSA determines at the decision block 306 that there is a virtual machine with a load below a threshold, the VSA, at block 308, selects a virtual machine from the set of virtual machines that the VSA identified as having a load below the threshold. The set of virtual machines may comprise the virtual machines identified at the decision block 306. To select the virtual machine from the set of identified virtual machines at the block 308, the VSA may use a round robin selection scheme, or any other type of load balancing selection scheme. In some cases, the virtual machine may be selected at random. In other cases, a weighted round robin algorithm or other weighted selection algorithm may be used to select the virtual machine from the set of virtual machines with a load below the threshold. For example, virtual machines may be weighted based on the capabilities of the computing system (e.g., client computing device 102 or server computing device 212) hosting the virtual machines. For instance, a host machine with more memory or a newer generation of processors may be weighted higher than a host machine with less memory or an older generation of processors such that the host machine with more memory or a newer generation of processors is selected more frequently than the host machine with less memory or older processors. As another example, a virtual machine may be selected based on the ability to access data or backup the virtual machine without a network. For instance, some virtual machines may be able to access data using a hotadd feature without accessing a network. The hotadd feature may enable accessing data or a new virtual disk without shutting down the virtual machine. Further, in some cases, the virtual machine may be located on the same datastore or repository as a virtual disk to be accessed via hotadd thereby enabling loading a new disk without accessing a network. In some cases, the virtual machine may access data using a storage area network (SAN) communicating over a fibre channel. In some embodiments, a virtual machine may be selected based on its ability to access data or be backed up over the SAN without accessing or transmitting the data over a LAN or other network that may be created between the VM host systems.

At block 310, the VSA assigns the job to the selected virtual machine that was selected at the block 308. Assigning the job to the selected virtual machine may include providing the job request to the selected virtual machine. Further, the block 310 may include updating one or more load tables associated with the virtual machine and/or set of virtual machines assigned to or allocated to the VSA. In addition, in some cases, the block 310 may include updating one or more load tables associated with one or more additional VSA associated with the virtual machine. For instance, in some cases a plurality of virtual server agents may be assigned a set of virtual machines. In such cases, the VSA may update each of the virtual server agents that have been assigned to the set of virtual machines. In some embodiments, the block 310 may include updating the VM management interface 245 with the assignment information. Advantageously, in certain embodiments, by updating the

VM management interface **245** with the assigned job information, a system or user (e.g., an administrator) may monitor the status of each of the VMs **204** and their assigned jobs or tasks by accessing a single system, the VM management interface **245**.

If the VSA determines at the decision block **306** that there are no virtual machines (e.g., a set of identified VMs is an empty set) assigned to the VSA with a load satisfying the threshold (e.g., at or below the threshold), the VSA may initiate creation of a new virtual machine at block **312**. Initiating creation of a new virtual machine may include selecting a virtual machine provider system (e.g., client computing device **102** or server computing device **212**) to host the new virtual machine. The virtual machine provider system may be selected based on the type of virtual machine to be created, the number of virtual machines already hosted by the virtual machine provider system, the type of job associated with the job request, or using any other factor or algorithm for selecting a virtual machine provider system. Further, in some cases, the virtual machine provider system may be selected using a load balancing scheme to ensure or reduce the probability that one virtual machine provider system is associated with a disproportionate load compared to other virtual machine provider systems in a group of virtual machine provider systems assigned to a set of VSAs. Once the virtual machine provider system has been selected, the VSA may use the virtual machine monitor **206** of the selected virtual machine provider system to create the new virtual machine.

After the new virtual machine has been created, the VSA may assign the job to the virtual machine at block **314**. Assigning the job to the new virtual machine may include providing the job request to the new virtual machine. In some cases, the block **314** may include one or more of the embodiments described with respect to the block **310**. For example, a load table associated with the newly created virtual machine may be updated to reflect the assignment of the job to the new virtual machine.

Example Virtual Machine Grouping Process

FIG. 4 illustrates a flowchart for an example virtual machine grouping process **400**. The process **400** can be implemented, at least in part, by any system that can group virtual machines based on virtual machine provider systems and/or virtual machine host systems. Advantageously, in certain embodiments, by grouping virtual machine provider systems, jobs may be allocated more efficiently to virtual machines based at least partially on the capabilities of the host systems. Further, management of the virtual machines and/or virtual machine provider systems may be simplified compared to information management systems with non-grouped virtual machine provider systems. Simplifying management can reduce management costs in terms of both time and money. In certain embodiments, the process **400** may be implemented, in whole or in part, by a storage manager **140**, a VSA coordinator **202**, and/or a VM management interface **245**, to name a few. Although any number of systems, in whole or in part, can implement the process **400**, to simplify discussion, portions of the process **400** will be described with reference to particular systems.

The process **400** begins at block **402** where, for example, a VSA coordinator **202** identifies a set of available virtual machine provider systems. These virtual machine provider systems may include any type of system that is capable of hosting a virtual machine **204**. For example, the virtual machine provider systems may include the client computing devices **102** and/or the server computing devices **212**. Further, the available virtual machine provider systems may

include any computing systems in the information management system **200** capable of hosting virtual machines. Alternatively, the available virtual machine provider systems may be limited to computing systems capable of hosting virtual machines that have been registered with the VSA coordinator **202** and/or the storage manager **140**.

At block **404**, the VSA coordinator **202** may access metadata associated with the set of available virtual machine provider systems identified at the block **402**. This metadata may include the identity of the features, specifications, and/or capabilities associated with each of the available virtual machine provider systems. For example, the metadata may include an amount of memory, a number of processors, or an amount of storage space for a particular virtual machine provider system, to name a few. As another example, the metadata may identify whether a particular virtual machine provider system is capable of supporting hot-add mode, or the ability to clone a virtual disk such that a virtual machine may continue running while the virtual disk is backed up to, for example, a secondary storage system. As yet another example, the metadata may include a geographic location of the virtual machine provider system, a network location of the virtual machine provider system, an age or amount of time in service of the virtual machine provider system, a security level associated with the virtual machine provider system, and any other information that may be used to categorize or group a virtual machine provider system with one or more additional virtual machine provider systems.

At block **406**, the VSA coordinator **202** groups the virtual machine provider systems into one or more categories based at least partially on the metadata accessed at the block **404**. In some embodiments, the VSA coordinator **202** may group the virtual machine provider systems based at least partially on grouping rules. These grouping rules may be accessed from the storage manager **140**. Alternatively, the grouping rules may be received from the virtual machine management interface **245**. Further, in some cases, the grouping rules may be generated by a user (e.g., an administrator). The grouping rules may include any type of rule for grouping virtual machine provider systems. For example, the grouping rules may be based on one or more types of metadata associated with the virtual machine provider systems.

The VSA coordinator **202**, at block **408**, assigns one or more virtual server agents **210** to each of the virtual machine provider system groups determined at the block **406**. Although in some cases a single virtual server agent may be assigned to a virtual machine provider system group, it is often advantageous to assign multiple virtual server agents. For one, operations may be completed faster because more streams may be maintained or initiated between the primary storage system **117** and the secondary storage system **118** when multiple virtual server agents are assigned to a group. An additional advantage, in some cases, is that each virtual server agent can serve as a failover virtual server agent for other virtual server agents assigned to the virtual machine provider system group. For instance, if a backup operation is in progress, and a virtual server agent fails, rather than the backup operation failing, the virtual machines assigned to the failed virtual server agent for backup may be reassigned to another virtual server agent in the group enabling the backup operation to proceed.

In some cases, at least some of the virtual server agents assigned to a virtual machine provider system group may include virtual server agents executing on the virtual machine provider systems of the virtual machine provider system group. Alternatively, or in addition, the virtual server

agents assigned to a virtual machine provider system group may be separate computing systems and/or may be executed on separate computing systems from the virtual machine provider systems included in the virtual machine provider system group.

In some embodiments, if an additional virtual machine provider system is added to the information management system **200**, the VSA coordinator **202** may access metadata for the new virtual machine provider system. The VSA coordinator **202** may then identify a group from the existing groups of virtual machine provider systems based at least partially on the metadata of the newly added virtual machine provider system. The VSA coordinator **202** may then assign the additional virtual machine provider system to the identified group. Assigning the additional virtual machine provider system to the identified group may include updating assignment information of the set of virtual server agents assigned to the identified group to include the additional virtual machine provider system. If a group cannot be identified based on the metadata of the newly added virtual machine, the VSA coordinator **202** may create a new group for the additional virtual machine provider system.

In certain implementations, the block **408** may include registering the virtual machine group with the VM management interface **245**. Further, in some cases, the VM management interface **245** may cause the VSA coordinator **202** to perform the process **400**.

A similar process to the process **400** may occur when the configuration of a virtual machine provider system changes. For example, if the VSA coordinator **202** detects that the configuration of a virtual machine provider system has changed, or if an administrator identifies to the VSA coordinator **202** to the configuration of a virtual machine provider system is changed, the VSA coordinator **202** may access the updated metadata for the reconfigured virtual machine provider system. Based on the updated metadata, the VSA coordinator **202** may reassign the virtual machine provider system to another group. Alternatively, if another group cannot be identified is associated with the updated metadata of the modified virtual machine provider system, the VSA coordinator **202** may create a new group or the modified virtual machine provider system. Further the VSA coordinator **202** may update assignment information for set of virtual server agents to include or to remove modified virtual machine provider system from the systems assigned to the set of virtual server agents.

In some embodiments, the process **400** may be performed periodically or may be performed each time a virtual machine provider system is added or removed from the information management system **200**. Further, in some cases, when a VSA is no longer assigned to a virtual machine group or when a virtual machine provider system group to which the VSA was assigned no longer includes any virtual machine provider systems and/or VMs, the VSA may be shut down or deactivated. Alternatively, the VSA may be reassigned to another virtual machine provider system group.

Example Virtual Server Agent Load Balancing Process

FIG. **5** illustrates a flowchart for an example virtual server agent load balancing process **500**. The process **500** can be implemented, at least in part, by any system that can balance the load of a set of virtual server agents. Advantageously, in certain embodiments, by balancing the load of the set of virtual server agents, operations may be completed faster and/or in a more efficient manner. In certain embodiments, the process **500** may be implemented, in whole or in part, by a storage manager **140**, a VSA **210**, a VM management

interface **245**, and a VSA coordinator **202**, to name a few. Although any number of systems, in whole or in part, can implement the process **500**, to simplify discussion, portions of the process **500** will be described with reference to particular systems.

It is of note that the process **500** may be used with respect to a number of types of operations including, for example, deduplication, file and/or virtual disk restoration, and file and/or virtual disk decryption/encryption, to name a few. Further, the process **500** may be used with respect to any type of operation that may involve communicating with the secondary storage subsystem **118**. However, to simplify discussion, and not to limit the process, the process **500** will be described primarily in the context of backup.

The process **500** begins at block **502** where, for example, a VSA coordinator **202** identifies a number of virtual machines to backup. The virtual machines may be identified in response to a backup command or a scheduled backup process. For example, a set of virtual machines may be scheduled for backup on a nightly, weekly, or monthly basis. Further, the set of virtual machines for backup may be identified based on a grouping of the virtual machines and/or a grouping of the virtual machine provider systems that host the set of virtual machines. In some cases, the set of virtual machines identified for backup may be identified based on the set of virtual server agents assigned to the virtual machines. In some embodiments, the virtual machines identified for backup may be identified by a user (e.g., an administrator) via, for example, the VM management interface **245**. In some implementations, virtual machines may be identified for backup based on one or more backup policies or other virtual machine management policies. These policies may be stored at or implemented by the storage manager **140**, a virtual machine management interface **245**, and/or a VSA coordinator **202**. Further, a determination of whether to implement one or more of the management policies may be made based on one or more factors associated with the virtual machines, such as instantiation time, time of last backup, type of VM, size of the VM, etc. In some cases, at least some of the factors may be based on the VSAs **210** associated with the VMs.

At block **504**, the VSA coordinator **202** may access metadata associated with the virtual machines identified at the block **502**. The metadata may be accessed from a table at a database or repository associated with the virtual machines and/or host systems of the virtual machines. In some cases, the VSA coordinator **202** may access the metadata by communicating with one or more virtual machine monitors **206** that facilitate management of the virtual machines **204** identified at the block **502**. Alternatively, or in addition, the VSA coordinator **202** may obtain access to the metadata by communicating with the one or more virtual server agents **210** assigned to the virtual machines identified at the block **502** and/or the host systems of the virtual machines.

The VSA coordinator **202**, at block **506**, identifies a number of virtual server agents **210** available to facilitate backup of the set of virtual machines identified at the block **502**. The set of virtual server agents **210** may be VSAs that have been assigned to the set of virtual machines and/or host systems of the virtual machines. For example, in reference to FIG. **2**, the virtual machines identified for backup at the block **502** include the virtual machines hosted by the server computing devices **212**, the set of VSAs identified at the block **506** may include one or more of the virtual server agents **210A** and virtual server agents **210C**.

At block **508**, the VSA coordinator **202** determines a number of streams available to each virtual server agent from the number of virtual server agents identified at the block **506**. The streams may be identified by accessing metadata associated with each of the identified virtual server agents. Further, the number of available streams may be based at least partially on a number of data paths and/or communication channels between the virtual server agents and one or more media agents of the secondary storage subsystem **118**. In some embodiments one or more of the blocks **506** and **508** may be optional. For example, in some cases the VSA coordinator **202** may determine a number of streams associated with a grouping of virtual machines and/or host computing systems of the virtual machines.

In some embodiments, the block **508** may include accessing and/or determining additional metadata associated with the virtual server agents identified at the block **506**. This metadata may include any type of information related to the load, capacity, or features of the virtual server agents. For example, the metadata may identify processing capacity or storage capacity of a virtual server agent, the number of processors, the type of processors, or the network bandwidth of each virtual server agent.

At block **510**, the VSA coordinator **202** distributes the virtual machines identified at the block **502** among the virtual server agents identified at the block **506** for backup to the secondary storage subsystem **118** based at least partially on the metadata associated with the virtual machines and the number of streams available to each virtual server agent. For example, the virtual machines may be distributed such that a virtual server agent with more available streams or a larger available bandwidth receives larger virtual machines while a virtual server agent with less available streams or a smaller available bandwidth may receive smaller virtual machines. As a second example, a virtual server agent with three streams may receive three virtual machines to backup, a virtual server agent with two streams may receive two virtual machines to backup, and a virtual server agent with one stream may receive one virtual machine to backup. In some embodiments, the block **510** may include distributing the virtual machine among the virtual server agents based partially on the metadata associated with the virtual server agents, such as the age of the virtual server agent or the geographic and/or network location of the virtual server agent.

In some cases, each of the virtual machines may be assigned to a virtual server agent for backup as part of a finite process. In other words, virtual machines may be queued for backup to a selected virtual server agent upon initiation of the backup operation or at the time of the scheduled backup. However, in other cases, virtual machines may be assigned to a virtual server agent on a rolling basis as a virtual server agent gains capacity to process the virtual machine, such as at the completion of the backup of another virtual machine. In other words, a virtual machine may only be assigned to the virtual server agent at a time that the virtual server agent is ready to begin processing the backup of the virtual machine. Advantageously, in certain embodiments, by assigning virtual machines to a virtual server agent for backup only when the virtual server agent is ready to begin backing up the virtual machine, load-balancing may be improved. For instance, in a case where there are ten virtual machines to backup and two virtual server agents, a round robin allocation performed at the start of backup would result in each virtual server agent being assigned five virtual machines to backup. In cases where the virtual machines are equal or close in size, five virtual machines per

virtual server agent may be optimal. However, in cases where the virtual machines differ substantially in size (e.g., by a factor 2 or 3), or the virtual server agents are unequal in capacity (e.g., network bandwidth or number of processors), such an even distribution may be suboptimal for completing backup in the shortest amount of time. Thus, in some cases, assigning a virtual machine for backup as a virtual server agent becomes available may be preferred.

In some cases the virtual machines may be assigned to the virtual server agents for backup using a round robin selection process. Alternatively, or in addition, virtual machines may be assigned to the virtual server agents for backup using a weighting process. For example, larger virtual machines, or virtual machines with a larger virtual disk, may be weighted higher than smaller virtual machines, virtual machines with a smaller virtual disk. In such cases, a virtual server agent that is assigned a larger virtual machine for backup may be assigned less virtual machines than a virtual server agent that is assigned smaller virtual machines.

Alternatively, or in addition, to the capabilities of the virtual server agents, a virtual server agent may be selected to backup a particular virtual machine based at least partially on the media agent to which the virtual server agent has an available stream. In certain embodiments, rather than distributing the virtual machines to the virtual server agents, the VSA coordinator **202** may distribute the virtual machines to streams without considering to which virtual server agent the stream belongs were to which media agent stream communicates with.

In some cases, backing up a virtual machine may include backing up files or data included in virtual storage of the virtual machine or a virtual disk. Alternatively, backing up the virtual machine may include backing up a file or a portion of a file that comprises the virtual machine.

TERMINOLOGY

Conditional language, such as, among others, “can,” “could,” “might,” or “may,” unless specifically stated otherwise, or otherwise understood within the context as used, is generally intended to convey that certain embodiments include, while other embodiments do not include, certain features, elements and/or steps. Thus, such conditional language is not generally intended to imply that features, elements and/or steps are in any way required for one or more embodiments or that one or more embodiments necessarily include logic for deciding, with or without user input or prompting, whether these features, elements and/or steps are included or are to be performed in any particular embodiment.

Unless the context clearly requires otherwise, throughout the description and the claims, the words “comprise,” “comprising,” and the like are to be construed in an inclusive sense, as opposed to an exclusive or exhaustive sense; that is to say, in the sense of “including, but not limited to.” As used herein, the terms “connected,” “coupled,” or any variant thereof means any connection or coupling, either direct or indirect, between two or more elements; the coupling or connection between the elements can be physical, logical, or a combination thereof. Additionally, the words “herein,” “above,” “below,” and words of similar import, when used in this application, refer to this application as a whole and not to any particular portions of this application. Where the context permits, words in the above Detailed Description using the singular or plural number may also include the plural or singular number respectively. The word “or” in reference to a list of two or more items, covers all of the

following interpretations of the word: any one of the items in the list, all of the items in the list, and any combination of the items in the list. Likewise the term “and/or” in reference to a list of two or more items, covers all of the following interpretations of the word: any one of the items in the list, all of the items in the list, and any combination of the items in the list.

Depending on the embodiment, certain acts, events, or functions of any of the algorithms described herein can be performed in a different sequence, can be added, merged, or left out altogether (e.g., not all described acts or events are necessary for the practice of the algorithms). Moreover, in certain embodiments, acts or events can be performed concurrently, e.g., through multi-threaded processing, interrupt processing, or multiple processors or processor cores or on other parallel architectures, rather than sequentially.

Systems and modules described herein may comprise software, firmware, hardware, or any combination(s) of software, firmware, or hardware suitable for the purposes described herein. Software and other modules may reside on servers, workstations, personal computers, computerized tablets, PDAs, and other devices suitable for the purposes described herein. Software and other modules may be accessible via local memory, via a network, via a browser, or via other means suitable for the purposes described herein. Data structures described herein may comprise computer files, variables, programming arrays, programming structures, or any electronic information storage schemes or methods, or any combinations thereof, suitable for the purposes described herein. User interface elements described herein may comprise elements from graphical user interfaces, command line interfaces, and other suitable interfaces.

Further, the processing of the various components of the illustrated systems can be distributed across multiple machines, networks, and other computing resources. In addition, two or more components of a system can be combined into fewer components. Various components of the illustrated systems can be implemented in one or more virtual machines, rather than in dedicated computer hardware systems. Likewise, the data repositories shown can represent physical and/or logical data storage, including, for example, storage area networks or other distributed storage systems. Moreover, in some embodiments the connections between the components shown represent possible paths of data flow, rather than actual connections between hardware. While some examples of possible connections are shown, any of the subset of the components shown can communicate with any other subset of components in various implementations.

Embodiments are also described above with reference to flow chart illustrations and/or block diagrams of methods, apparatus (systems) and computer program products. Each block of the flow chart illustrations and/or block diagrams, and combinations of blocks in the flow chart illustrations and/or block diagrams, may be implemented by computer program instructions. Such instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the acts specified in the flow chart and/or block diagram block or blocks.

These computer program instructions may also be stored in a non-transitory computer-readable memory that can direct a computer or other programmable data processing apparatus to operate in a particular manner, such that the

instructions stored in the computer-readable memory produce an article of manufacture including instruction means which implement the acts specified in the flow chart and/or block diagram block or blocks. The computer program instructions may also be loaded onto a computer or other programmable data processing apparatus to cause a series of operations to be performed on the computer or other programmable apparatus to produce a computer implemented process such that the instructions which execute on the computer or other programmable apparatus provide steps for implementing the acts specified in the flow chart and/or block diagram block or blocks.

Any patents and applications and other references noted above, including any that may be listed in accompanying filing papers, are incorporated herein by reference. Aspects of the invention can be modified, if necessary, to employ the systems, functions, and concepts of the various references described above to provide yet further implementations of the invention.

These and other changes can be made to the invention in light of the above Detailed Description. While the above description describes certain examples of the invention, and describes the best mode contemplated, no matter how detailed the above appears in text, the invention can be practiced in many ways. Details of the system may vary considerably in its specific implementation, while still being encompassed by the invention disclosed herein. As noted above, particular terminology used when describing certain features or aspects of the invention should not be taken to imply that the terminology is being redefined herein to be restricted to any specific characteristics, features, or aspects of the invention with which that terminology is associated. In general, the terms used in the following claims should not be construed to limit the invention to the specific examples disclosed in the specification, unless the above Detailed Description section explicitly defines such terms. Accordingly, the actual scope of the invention encompasses not only the disclosed examples, but also all equivalent ways of practicing or implementing the invention under the claims.

To reduce the number of claims, certain aspects of the invention are presented below in certain claim forms, but the applicant contemplates the various aspects of the invention in any number of claim forms. For example, while only one aspect of the invention is recited as a means-plus-function claim under 35 U.S.C. sec. 112(f) (AIA), other aspects may likewise be embodied as a means-plus-function claim, or in other forms, such as being embodied in a computer-readable medium. Any claims intended to be treated under 35 U.S.C. § 112(f) will begin with the words “means for”, but use of the term “for” in any other context is not intended to invoke treatment under 35 U.S.C. § 112(f). Accordingly, the applicant reserves the right to pursue additional claims after filing this application, in either this application or in a continuing application.

What is claimed is:

1. A computer-readable storage medium storing instructions that when executed by a computer cause the computer to perform a method for reducing virtual machine (VM) proliferation, the method comprising: receiving a job request at a computing device comprising computer hardware; determining an expected load for the job request; selecting a threshold load based at least partially on the expected load for the job request; determining a load for each VM within a set of VMs; determining whether a load of at least one VM is below the threshold load; and

59

in response to determining that the load of at least one VM is below the threshold load, selecting a VM from a subset of VMs, wherein each VM in the subset of VMs is associated with a load that is below the threshold load; and

assigning a job associated with the job request to the selected VM.

2. The computer-readable storage medium of claim 1, wherein determining the load for each VM from the set of VMs comprises accessing one or more VM monitors to obtain the load for each VM.

3. The computer-readable storage medium of claim 1, wherein the job request is a data protection operation according to a storage policy, wherein the storage policy is a collection of settings or preferences for performing data protection operations on data associated with the storage policy.

4. The computer-readable storage medium of claim 1, wherein the job request is received from a virtual server agent coordinator.

5. The computer-readable storage medium of claim 1, wherein a virtual server agent is configured to at least partially manage the set of VMs.

6. The computer-readable storage medium of claim 5, wherein the virtual server agent is part of a primary storage system.

7. The computer-readable storage medium of claim 5, wherein the virtual server agent is configured to communicate with one or more media agents of a secondary storage system via one or more streams between the virtual server agent and each media agent of the one or more media agents.

8. The computer-readable storage medium of claim 1, wherein selecting the VM comprises selecting a VM based at least partially on a load balancing scheme.

9. The computer-readable storage medium of claim 8, wherein the load balancing scheme is specified by a storage policy associated with data to be modified or copied by the job request,

wherein the storage policy is a collection of settings or preferences for performing data protection operations on data associated with the storage policy.

10. The computer-readable storage medium of claim 1, the method further comprising:

identifying a job type of the job request; and selecting the threshold load based at least partially on the job type.

11. The computer-readable storage medium of claim 1, the method further comprising:

initiating creation of a new VM, wherein the initiating of creation of the new VM comprises:

selecting a computing system from a set of computing systems, the set of computing systems hosting the set of VMs; and

causing a VM monitor of the selected computing system to create the new VM.

12. The computer-readable storage medium of claim 11, wherein selecting the computing system is based at least partially on a job type of the job request.

13. The computer-readable storage medium of claim 11, wherein selecting the computing system comprises selecting the computing system based at least partially on a load balancing scheme.

14. A system for reducing VM proliferation, the system comprising:

one or more computing devices comprising computer hardware, the one or more computing devices configured to:

60

receive a job request to perform a data protection operation according to a storage policy,

wherein the storage policy is a collection of settings or preferences for performing data protection operations on data associated with the storage policy;

determine an expected load for the job request;

select a threshold load based at least partially on the expected load for the job request;

access load information for each VM from a set of available VMs;

identify, based at least partially on the load information for each VM, a subset of VMs from the set of available VMs,

wherein each VM in the subset of VMs is associated with a load that is below the threshold load; and

select a VM from the subset of VMs and assign a job associated with the job request to the selected VM.

15. The system of claim 14, wherein the data protection operation relates to pruning or aging data.

16. The system of claim 14, wherein the data protection operation is a backup operation for backing up primary data to a secondary storage system.

17. The system of claim 14, wherein the one or more computing devices are further configured to initiate creation of a new VM and to assign the job associated with the job request to the new VM.

18. The system of claim 17, wherein the one or more computing devices are further configured to initiate the creation of the new VM by:

selecting a computing system from a set of computing systems, the set of computing systems configured to host a set of VMs; and

causing a VM monitor of the selected computing system to create the new VM.

19. The system of claim 17, wherein the one or more computing devices are further configured to provide information associated with the new VM to a VM management interface configured to present VM status information to a user.

20. The system of claim 14, wherein the one or more computing devices are further configured to access one or more VM monitors to obtain the load information for each VM.

21. The system of claim 14, wherein the one or more computing devices are further configured to communicate with one or more media agents of a secondary storage system via one or more streams.

22. A system for reducing VM proliferation, the system comprising:

one or more computing devices comprising computer hardware, the one or more computing devices configured to:

receive a job request;

determine an expected load for the job request;

select a threshold load based at least partially on the expected load for the job request;

access load information for each VM from a set of available VMs;

identify, based at least partially on the load information for each VM, a subset of VMs from the set of available VMs; wherein each VM in the subset of VMs is associated with a load that is below the threshold load; and

and

61

create a new VM when the subset of VMs is a zero subset;
and assign the job associated with the job request to the
new VM.

* * * * *

62